

To complete this assignment follow the setup instructions from the course Web page (if not already followed). The setup instructions bring you to the point where you can compile the *cpu-only* examples but for this assignment that won't be necessary. Then follow the instructions in the first problem to complete and test the setup.

**Problem 0:** Check out the latest copy of the balls project base code and the include directory. The transcript below shows the steps needed.

(a) Read the following overview of how the code operates.

This code includes something like a solution to Homework 3, except that direct cuda calls were not made and the code returns the actual proximity pairs, not just a count. The code computing proximity pairs runs inefficiently due to data access problems and other issues, the point of this homework assignment is to improve it.

Here is an outline of how the affected code operates:

The code in `balls.cc:cuda_contact_pairs_find` sorts the balls (`phys`) in `z` order using class `PSList`. An array on cuda, `z_sort_indices`, maps a `z`-order index to a balls array index. For example, `z_sort_index[5] == 12` means that ball number 12 (the index to `cpu array phys` and `gpu arrays ball_x`) has the sixth smallest `z` value. This array is initialized in `balls.cc:cuda_contact_pairs_find` and sent over to CUDA. The `z` max values are also placed in an array and sent to CUDA.

The routine launches a CUDA kernel, `pass_sched`, to compute the proximity pairs, and reads back this data, in CUDA array `cuda_prox`. The data in `cuda_prox` is unpacked in routine `contact_pairs_find`.

The problems below are related to the data moving between CPU and GPU, and data moving between the different memory areas in the GPU. Here is a summary of data movement that the proximity code currently induces: The kernel `pass_sched` reads the `z_sort_indices` and `z` max arrays, and also reads the data in `balls_x` that is already present on the GPU. (That is, the data for `balls_x` is only sent to the GPU when balls are added or removed. Otherwise the GPU keeps using the copy it has.) The kernel writes only the array `cuda_prox`.

Each thread of the kernel, `pass_sched` in file `balls-kernel.cu`, performs one outer iteration of the `idx9` loop from `contact_pairs_proximity_check`. That is, each thread has its own `idx9` value and must iterate `idx1` values until there is no possibility of proximity.

Unlike `contact_pairs_proximity_check`, `pass_sched` does not check for proximity between tiles and balls, only balls and balls. (This was to make the assignment less tedious.) The CPU code, in `contact_pairs_proximity_check`, prepares a `Contact` structure for each nearby `idx1` it finds. The corresponding code in `pass_sched` simply passes the values `idx1s` back to the CPU code. It does this by compressing `idx1` values to 8 bits (by subtracting them from `idx9`) and backs them in to a 64 bit integer names `offsets`. The CPU code uses `offsets` to create the contact structures.

See the comments in the code for additional information.

(b) Familiarize yourself with the relevant UI:

Various features can be toggled on and off, and their state is shown in green at the top of the display. An indicator will blink if the corresponding should be turned off when measuring timing. This includes CPU physics and correctness verification.

Press 'F4' to toggle CUDA proximity checking.

Press 'v' to toggle verification. If proximities are incorrect execution will exit with an abort. Within a debugger you can issue the command "frame 2" (or sometimes "frame 3") to show where

the verification failed. There will be a statement like `ASSERTS( something )` where something is false but should be true. Verification is **time consuming** and so there will be no speedup as long as its on. That's why the indicator blinks when it is on.

Press 't' to run a benchmark that drops thousands of balls on the platform with certain constants, such as friction set to special values. (Edit `benchmark_setup` to change the behavior.)

Press 'T' to run the benchmark with fewer balls, this is helpful when on a slower system or testing slower code.

Press 'a' to switch between CPU and GPU physics.

Press 'c' to toggle between natural colors and color based on CUDA block assignment. (See VAR 'Color by Block in Pass'.)

Press 'q' to toggle the value of variable `World::opt_debug` between true and false. The current value is the first number after "Debug Options." The variable `World::opt_debug` currently does nothing, it can be used to test ideas for your solution. Variable `opt_debug` is also available in CUDA.

Press 'Q' to toggle the value of variable `World::opt_debug2`. This also does nothing and is available for working on the assignment. Variable `opt_debug2` is also available in CUDA.

Timings for the scheduling pass are shown on the upper right (it will be necessary to make the window wider). CUDA Prox shows the time in the routine `World::cuda_contact_pairs_find()`, Sched shows the total time for scheduling, CPU Prox shows the time in `World::contact_pairs_find()`. When there are many balls (press t) and the CPU is computing proximity the value of CPU Prox should be a large fraction of Sched. With CUDA computing proximity CPU Prox will be much smaller, it's not zero because the CPU must still compute tile interactions and it must unpack the CUDA data. These timings are provided by `frame_timer.user_timer_X` calls, feel free to add your own calls to time other parts of the code.

The F2010 HW4 line shows whether CUDA proximity testing is one. The CPU test/ball shows how many proximity tests per ball are performed; total shows the total number of tests.

(c) Check out the code using the following transcript. Don't forget to update the include directory.

```
# Transcript of commands to check out and run code for
# this assignment.

[orion.ece.lsu.edu] % cd ~
/home/faculty/koppel

# Check out a copy of the balls code into a directory named hw4.
[orion.ece.lsu.edu] % svn co https://svn.ece.lsu.edu/svn/gp/proj-base/balls hw4
A   hw4/balls-shdr.cc
A   hw4/balls.png
# ...
Checked out revision 242.

# Check out or update the include directory.

# Because the repository include directory has changed this step needs
# to be repeated unless the last checkout was very recent.
# (An easier way to update is to execute "svn update" in the include
# directory.)
[orion.ece.lsu.edu] % svn co https://svn.ece.lsu.edu/svn/gp/include
A   include/gl-buffer.h
```

```

A    include/misc.h
A    include/texture-util.h
A    include/util.h
A    include/gp
A    include/glexthfuncs.h
A    include/coord.h
A    include/pstring.h
A    include/shader.h
A    include/cuda-util.h
Checked out revision 242.

```

# Go to the balls directory (named hw4), compile, and run the code.

```

[orion.ece.lsu.edu] % cd hw4
/home/faculty/koppel/example/hw4
[orion.ece.lsu.edu] % make -j 4
/usr/local/cuda/bin/nvcc -Xcompiler -Wall --ptxas-options=-v -I/usr/X11R6/include/
-I/opt/local/include -I../include -I../..//include -g -Xcompiler -Wno-strict-aliasing
-Xcompiler -Wno-parentheses -c balls-kernel.cu
# ...

```

```

[orion.ece.lsu.edu] % balls
S GL_VENDOR: "NVIDIA Corporation"
S GL_RENDERER: "Quadro FX 3800/PCI/SSE2"
S GL_VERSION: "3.2.0 NVIDIA 190.53"

```

(d) Check out, build, and play a bit with the code.

**Problem 1:** In this problem determine various limits on how fast contact pairs can be computed. The limits will be measured in proximity tests per second, and called an execution rate. (A proximity test is what is done by an iteration of the `idx1` loop in `contact_pairs_proximity_check`.) Let  $n$  denote the number of balls and let  $p$  denote the average number of proximity tests per ball. The total work done by the code is then  $np$  proximity tests. Call the rate of proximity tests, measured in proximity tests per second, the *execution rate*.

Base your answers to the questions below on the GeForce 8800 GTX as described in the 2006 NVIDIA technical brief and on the code checked out above.

(a) Examine the PTX code for `pass_sched` and estimate a number of instructions in a `idx1` loop iteration. Choose a useful iteration or iterations. Provide the number of instructions and give the execution rate based on this number of instructions assuming a sufficiently large number of threads and unlimited memory bandwidth. Bonus points for usefully taking into account the fraction of proximity tests for which `idx1` is nearby.

(b) Compute the total amount of data sent from the CPU to the GPU for  $n$  balls.

(c) Compute the execution rate based on the amount of data fetched from device memory to the multiprocessors. If a data item is fetched twice, count it twice, and if a memory fetch has a width of 32 bytes but only four bytes are used, count that as 32 bytes. See the CUDA Programming guide. Base the number of fetches on a machine with compute capability 1.3, but base the timing and bandwidth on the 8800. To get the execution rate first compute the total amount of data fetched

for  $n$  balls and  $p$  test per ball, then divide by the maximum data transfer rate (based on the 8800 description), then divide by  $np$ .

(d) Compute the execution rate based upon the amount of data fetched, but assuming perfect coalescing and alignment and assuming unlimited multiprocessor cache (count a data item once, even if it was used multiple times).

**Problem 2:** Improve performance of the code by using the texture cache. *Hint: This should only take a few lines of code.*

(a) Modify the code so that it uses the texture cache.

(b) Measure the change in performance and comment on it.

**Problem 3:** Show how shared memory can be used to improve performance. Outline how shared memory would be used, code is optional.