


**CSCI-GA.3033-004**  
**Graphics Processing Units (GPUs): Architecture and Programming**  
**Programming Assignment 1**

In your first programming assignment, you will implement a GPU friendly application and see yourself when a GPU friendly application will yield better performance than sequential version and when not.

You will write CUDA program to determine the heat distribution in a space using synchronous iteration on a GPU. We have 2-dimensional square space and simple boundary conditions (edge points have fixed temperatures). The objective is to find the temperature distribution within. The temperature of the interior depends upon the temperatures around it. We find the temperature distribution by dividing the area into a fine mesh of points,  $h_{i,j}$ . The temperature at an inside point is calculated as the average of the temperatures of the four neighboring points, as shown in the following figure.



The edge points are when  $i = 0$ ,  $i = n-1$ ,  $j = 0$ , or  $j = n-1$ , and have fixed values corresponding to the fixed temperatures of the edges. The temperature at each interior point is calculated as:

$$h_{i,j} = \frac{h_{i-1,j} + h_{i+1,j} + h_{i,j-1} + h_{i,j+1}}{4}$$

( $0 < i < n$  ;  $0 < j < n$ , remember that edge points have fixed values)

Stopping condition: You can have a fixed number of iterations or when the difference between two consecutive iterations (calculated as average among all points) is less than or equal to some number  $e$ . For this assignment, we will use fixed number of iterations.

Assume we have  $(n \times n)$  points (including edge points). The edge points are fixed at 80F, except points  $(0, 10)$  to  $(0, 30)$  inclusive have temperature of 150F. Assume also that all internal points are initialized to zero.

### What you have to do:

1. [10 points] Write CUDA version with all optimizations you can think of, in the provided file `heatdist.cu`. The simplest way to compile your code is:  
**`nvcc -o heatdist heatdist.cu`**  
The generated binary will have the name: *heatdist*.  
Type: **`./heatdist`** and press enter to see the usage. Then start implementing your GPU version in the designated place in the `*.cu` file. Read the comments on that file before starting to code.
2. Number of iterations is fixed to 50. Starting with  $n = 100$  and then doubling in each experiment (i.e. 100, 200, 400, ...), record the time taken by the CPU and that taken by the GPU. The provided file already calculates the time for you and print it on the screen,
3. [3 points] Draw a bar-graph showing  $n$  (x-axis) and the time (y-axis) for the 2 versions of the program (CPU and GPU) (2 bars per  $n$  on the same graph).
4. [24 points] What are your conclusions regarding:
  - When is GPU usage more beneficial (at which  $n$ )? and why?
  - When is the speedup (i.e. time of CPU version / time of CUDA version) at its lowest? And why?
  - When is the speedup at its highest? And why?
5. [13 points] Design your own experiments and determine whether the number of iterations has an effect on GPU performance. If yes, what is it? and why? If not, why not? Describe your experiment(s) and explain your conclusions.

### What to submit:

- The source code `heatdist.cu`
- A report that contains the graphs (#3) and answers to the questions mentioned above (#4 and #5).

Put all the above in a zip file named by your **lastname-firstname.zip** and email it to the grader.

I hope that after implementing this programming assignment, **you learned the following items** first-hand:

- How to allocate/free memory space at the device.
- How to move data from host to device, and back, if needed.
- How to launch a kernel from host to device.
- Get a feel of when the problem size is enough to get better performance on device when the kernel is GPU friendly.

So .....

# Have Fun!