



# GUJARAT TECHNOLOGICAL UNIVERSITY

**Master of Engineering**  
**Subject Code: 3720221**  
**Semester – II**  
**Subject Name: GPU Computing**

**Type of course:** Elective

**Prerequisite:**

**Rationale:** Almost, all modern processors are multi-core, and in order to make effective use of them, you need to run multiple threads. When accelerators, such as GPUs, are available, then the number of threads needs to be even larger. This course teaches you how to organize the computations of the threads so that they work together and performs the required computations efficiently, making good use of the available hardware resources. We will focus on using GPUs for general purpose computing, rather than for graphics.

**Teaching and Examination Scheme:**

Teaching Scheme			Credits C	Examination Marks				Total Marks
L	T	P		Theory Marks		Practical Marks		
				ESE (E)	PA (M)	ESE (V)	PA (I)	
3	0	2	4	70	30	30	20	150

**Content:**

Sr. No.	Content	Total Hrs	% Weightage
1	<b>Introduction:</b> History, GPU Architecture, Clock speeds, CPU / GPU comparisons, Heterogeneity, Accelerators, Parallel Programming, CUDA OpenCL / OpenACC, Kernels Launch parameters, Thread hierarchy, Warps/Wavefronts, Threadblocks/Workgroups, Streaming multiprocessors, 1D/2D/3D thread mapping, Device properties, Simple Programs	13	27
2	<b>Memory:</b> Memory hierarchy, DRAM / global, local / shared, private / local, textures, Constant Memory, Pointers, Parameter Passing, Arrays and dynamic Memory, Multi-dimensional Arrays, Memory Allocation, Memory copying across devices, Programs with matrices, Performance evaluation with different memories	7	15
3	<b>Synchronization:</b> Memory Consistency, Barriers (local versus global), Atomics, Memory fence. Prefix sum, Reduction. Programs for concurrent Data Structures such as Worklists, Linked-lists. Synchronization across CPU and GPU <b>Functions:</b> Device functions, Host functions, Kernels functions, Using libraries (such as Thrust), and developing libraries.	10	21
4	<b>Support:</b> Debugging GPU Programs. Profiling, Profile tools, Performance aspects <b>Streams:</b> Asynchronous processing, tasks, Task-dependence, Overlapped data transfers, Default Stream, Synchronization with streams. Events, Event-based-Synchronization - Overlapping data transfer and kernel execution, pitfalls.	8	17



# GUJARAT TECHNOLOGICAL UNIVERSITY

## Master of Engineering

Subject Code: 3720221

5	<b>Case Studies:</b> Image Processing, Graph algorithms, Simulations, Deep Learning	5	10
6	<b>Advanced topics:</b> Dynamic parallelism, Unified Virtual Memory, Multi-GPU processing, Peer access, Heterogeneous processing	5	10
	<b>Total</b>	48	<b>100%</b>

### Reference Books:

1. David Kirk and Wen-mei Hwu, Programming Massively Parallel Processors: A Hands-On Approach, 2nd Edition, Publisher: Morgan Kaufman, 2012, ISBN: 9780124159921.
2. Shane Cook, CUDA Programming: A Developer's Guide to Parallel Computing with GPUs, Morgan Kaufman; 2012 (ISBN: 978-0124159334)

### Course Outcomes:

After learning the course the students should be able to:

Sr. No.	CO statement	Marks % weightage
CO-1	Define terminology commonly used in parallel computing, such as efficiency and speedup.	10%
CO-2	Describe common GPU architectures and programming models.	20%
CO-3	Implement efficient algorithms for common application kernels, such as matrix multiplication.	20%
CO-4	Given a problem, develop an efficient parallel algorithm to solve it.	25%
CO-5	Given a problem, implement an efficient and correct code to solve it, analyze its performance, and give convincing written and oral presentations explaining the achievements.	25%

### Distribution of marks weightage for cognitive level

Bloom's Taxonomy for Cognitive Domain	Marks % weightage
Recall	5
Comprehension	10
Application	20
Analysis	25
Evaluate	25
Create	15

### Practical List:

1. Write CUDA code to compute the squares of the first N integers.
2. Write CUDA code to determine the following:
  - i. data transfer bandwidth from host to device



# GUJARAT TECHNOLOGICAL UNIVERSITY

## Master of Engineering

**Subject Code: 3720221**

- ii. data transfer bandwidth from device to host
  - iii. data transfer bandwidth from host to device using pinned memory
  - iv. data transfer bandwidth from device to host using pinned memory
  - v. kernel creation overhead
3. Implement matrix multiplication on the CPU and GPU(without using shared memory), and compare their relative performances in terms of GFlop/s and report your performance results.
  4. Write code for matrix multiplication using shared memory and compare its performance with CPU code. Determine the best tile-size for matrix multiplication using shared memory for your GPU machine.
  5. Implement the 1-D convolution kernel and compare the performance with and without shared memory.
  6. Perform 2-D convolution on an  $n \times n$  matrix with an  $m \times m$  mask and determine the number of halo cells required for it.
  7. Implement 2-D convolution with data in shared memory. Also, analyze the reduction in bandwidth from use of shared memory in 2-D convolution.
  8. Write code for histogramming without atomic operations and experiment with it to find out the fraction of times that it gives incorrect results and report it.
  9. Rewrite program 8 so that it uses atomic operations, and compare its performance with it and report it.
  10. Write code to perform reduction using atomic operations on global memory. (ii) Rewrite this code to use shared memory in order to reduce the overhead of atomic operations. (iii) Compare the performances of the above two with the reduction algorithm we had discussed earlier.
  11. Write code with different threads performing atomic operations on different, but adjacent, locations in memory. Compare the performance with different data sizes.