

# **GPU Programming in MATLAB**

Nikolaos Ploskas

Nikolaos Samaras



# Contents

Chapter 1. Introduction	1
1.1. Chapter Objectives	1
1.2. Parallel Programming	1
1.3. GPU Programming	6
1.4. CUDA Architecture	7
1.5. Why GPU Programming in MATLAB? When to Use GPU Programming?	10
1.6. Our Approach: Organization of the Book	15
1.7. Chapter Review	16
Chapter 2. Getting Started	19
2.1. Chapter Objectives	19
2.2. Hardware Requirements	19
2.3. Software Requirements	21
2.4. Chapter Review	31
Chapter 3. Parallel Computing Toolbox	33
3.1. Chapter Objectives	33
3.2. Product Description & Objectives	33
3.3. Parallel for-Loops (parfor)	35
3.4. Single Program Multiple Data (spmd)	48
3.5. Distributed & Codistributed Arrays	53
3.6. Interactive Parallel Development (pmode)	59
3.7. GPU Computing	61
3.8. Clusters & Job Scheduling	62
3.9. Chapter Review	66
Chapter 4. Introduction to GPU Programming in MATLAB	67
4.1. Chapter Objectives	67
4.2. GPU Programming Features in MATLAB	67
4.3. GPU Arrays	68
4.4. Built-in MATLAB Functions for GPUs	75
4.5. Element-Wise MATLAB Code on GPUs	88
4.6. Chapter Review	101
Chapter 5. GPU Programming on MATLAB toolboxes	103
5.1. Chapter Objectives	103
5.2. Communications System Toolbox	103
5.3. Image Processing Toolbox	122
5.4. Neural Network Toolbox	127

5.5. Phased Array System Toolbox	148
5.6. Signal Processing Toolbox	154
5.7. Statistics and Machine Learning Toolbox	156
5.8. Chapter Review	162
Chapter 6. Multiple GPUs	163
6.1. Chapter Objectives	163
6.2. Identify & Run Code on a Specific GPU Device	163
6.3. Examples Using Multiple GPUs	170
6.4. Chapter Review	189
Chapter 7. Run CUDA or PTX Code	191
7.1. Chapter Objectives	191
7.2. A Brief Introduction to CUDA C	191
7.3. Steps to Run CUDA or PTX Code on a GPU through MATLAB	195
7.4. Example: Vector Addition	203
7.5. Example: Matrix Multiplication	206
7.6. Chapter Review	209
Chapter 8. MATLAB MEX Functions Containing CUDA Code	211
8.1. Chapter Objectives	211
8.2. A Brief Introduction to MATLAB MEX Files	211
8.3. Steps to Run MATLAB MEX Functions on GPU	215
8.4. Example: Vector Addition	224
8.5. Example: Matrix Multiplication	227
8.6. Chapter Review	231
Chapter 9. CUDA-Accelerated Libraries	233
9.1. Chapter Objectives	233
9.2. Introduction	233
9.3. cuBLAS	234
9.4. cuFFT	237
9.5. cuRAND	241
9.6. cuSOLVER	245
9.7. cuSPARSE	248
9.8. NPP	252
9.9. Thrust	257
9.10. Chapter Review	259
Chapter 10. Profiling Code & Improving GPU Performance	261
10.1. Chapter Objectives	261
10.2. MATLAB Profiling	261
10.3. CUDA Profiling	274
10.4. Best Practices for Improving GPU Performance	277
10.5. Chapter Review	284
References	285
List of Examples	287
Index	291