

GPU-based ultra-fast dose calculation using a finite size pencil beam model

Xuejun Gu¹, Dongju Choi², Chunhua Men¹, Hubert Pan¹,
Amitava Majumdar² and Steve B Jiang¹

¹ Department of Radiation Oncology, University of California San Diego, La Jolla, CA 92037-0843, USA

² San Diego Supercomputer Center, University of California San Diego, La Jolla, CA 92093, USA

E-mail: sbjiang@ucsd.edu

Received 10 August 2009, in final form 15 September 2009

Published 1 October 2009

Online at stacks.iop.org/PMB/54/6287

Abstract

Online adaptive radiation therapy (ART) is an attractive concept that promises the ability to deliver an optimal treatment in response to the inter-fraction variability in patient anatomy. However, it has yet to be realized due to technical limitations. Fast dose deposit coefficient calculation is a critical component of the online planning process that is required for plan optimization of intensity-modulated radiation therapy (IMRT). Computer graphics processing units (GPUs) are well suited to provide the requisite fast performance for the data-parallel nature of dose calculation. In this work, we develop a dose calculation engine based on a finite-size pencil beam (FSPB) algorithm and a GPU parallel computing framework. The developed framework can accommodate any FSPB model. We test our implementation in the case of a water phantom and the case of a prostate cancer patient with varying beamlet and voxel sizes. All testing scenarios achieved speedup ranging from 200 to 400 times when using a NVIDIA Tesla C1060 card in comparison with a 2.27 GHz Intel Xeon CPU. The computational time for calculating dose deposition coefficients for a nine-field prostate IMRT plan with this new framework is less than 1 s. This indicates that the GPU-based FSPB algorithm is well suited for online re-planning for adaptive radiotherapy.

1. Introduction

Intensity-modulated radiation therapy (IMRT) is capable of delivering a highly conformal radiation dose to a complex static target volume. However, due to inter-fraction variation of patient anatomy, an optimal IMRT plan designed before the treatment may become less optimal or even totally unacceptable at some point during the treatment course (Yan *et al*

2005). With the use of on-board volumetric imaging techniques, this variation can be readily measured before each treatment fraction and then utilized to make useful modification of the original treatment plan (Yan *et al* 1997, Jaffray *et al* 2002). This procedure is often called adaptive radiation therapy (ART) (Yan *et al* 1997, Wu *et al* 2002, 2004, 2008, Birkner *et al* 2003, Mohan *et al* 2005, de la Zerda *et al* 2007, Lu *et al* 2008, Fu *et al* 2009, Godley *et al* 2009). One way to implement ART is based on offline re-planning during the treatment course. The schedule for imaging and re-planning is designed based on the modeling of the systematic and random components of the inter-fraction variations. An alternative implementation of ART is based on online re-planning, where a new plan is developed immediately after acquiring the patient's anatomy while the patient is still lying on the treatment couch. The online approach is preferable, as it allows real-time adaptations of the treatment plan to daily anatomical variations. However, it is challenging to implement online ART in clinical practice due to various technical barriers. One such barrier is the requirement of a very fast treatment planning process that can be completed within a few minutes. This constraint is extremely difficult, if not impossible, to satisfy with the central processing unit (CPU)-based computational framework that is currently used in clinical settings.

A typical online adaptive re-planning process includes the following main computational tasks: (1) the reconstruction of in-room computed tomography (CT) images, (2) the segmentation of the target and organs at risk (OARs) in CT images, (3) the computation of the dose distribution on the patient's new geometry and (4) the optimization of a new treatment plan. Each of these tasks takes a relatively long time under the current CPU-based treatment planning framework and therefore makes online re-planning impractical.

Traditionally, high-speed computation overwhelmingly relies on the advancement of the processing speed of the CPU. However, the uncertainty of the continuation of exponential growth in processing power as described by Moore's law (Dubash 2005) has led to renewed interest in multi-core and many-core computational architectures. CPU-clustered traditional supercomputers have been used for solving computationally intensive problems for decades. Despite their great computational power, traditional supercomputers are neither readily available nor accessible to most clinical users due to the prohibitively high cost of facility deployment and maintenance.

The concept of using stream processors in general-purpose graphic processing units (GPUs) depicts an innovative scenario of handling massive floating-point computation and making high-performance computation affordable to general users. GPUs are especially well suited for problems that can be expressed as data-parallel computations, such as programs of high arithmetic intensity (the ratio of arithmetic operations to memory operations) which need to be executed on multiple data elements in parallel (NVIDIA 2009). With affordable graphic cards such as NVIDIA's GeForce, GTX and Tesla series, GPU-based computing has recently been utilized to speed up computational tasks relevant to radiotherapy, such as CBCT reconstruction, deformable image registration and dose calculation (Sharp *et al* 2007, Meihua *et al* 2007, Preis *et al* 2009, Riabkov *et al* 2008, Samant *et al* 2008, Xu and Mueller 2007, Hissoiny *et al* 2009, Noe *et al* 2008, Jacques *et al* 2008). Among them, Jacques *et al* (2008) and Hissoiny *et al* (2009) have explored GPUs for fast dose computation. These two groups focused on the acceleration of the superposition/convolution algorithm. Jacques *et al* (2008) used a combination of a Digital Mars' D program and the Compute Unified Device Architecture (CUDA) development environment. Hissoiny *et al* (2009) modified part of a public domain treatment planning system (PlanUNC) using CUDA and achieved 10–15 times speedup. A larger speedup factor is expected for dose calculation implementations that are designed specifically for GPUs.

By exploring the massive computational resource of GPUs, we are now offered the prospect of overcoming the computational bottleneck of real-time online re-planning. At the University of California San Diego (UCSD), we are developing a supercomputing online re-planning environment (SCORE) based on reasonably priced and readily available GPUs. As part of this effort, we report in this paper the development of a CUDA-based parallel computing framework for ultra-fast dose calculation using a finite-size pencil beam (FSPB) model. The remainder of this paper is organized as follows. In section 2.1, the general FSPB model will be introduced and a specific FSPB kernel is described. The CUDA implementation of a general FSPB framework is detailed in section 2.2. Section 3 presents experimental results of dose calculation in the case of a water phantom and the case of a prostate cancer patient. The computational time is compared between GPU and CPU implementations. Finally, conclusions and discussion are provided in section 4.

2. Methods and materials

2.1. FSPB model

In the process of IMRT inverse planning, a broad beam is divided into small rectangular or square beamlets, and the contribution of each beamlet to every relevant voxel (often called dose deposition coefficient or D_{ij}) is calculated using a dose engine. FSPB models are particularly well suited for such calculations since a beamlet is naturally a finite-size pencil beam, and the accuracy of the models is sufficient in most clinical situations (Bourland and Chaney 1992, Ostapiak *et al* 1997, Jiang 1998, Jelen *et al* 2005). The major assumptions in an FSPB model include (1) the broad beam from a point source can be divided into identical beamlets and (2) the dose to a point is the integration of the contribution dose from all beamlets to that point. Mathematically, the dose deposited at a point $P(x, y, z)$ can be written as

$$D(x, y, z) = \sum_i w_i D_i^p(x, y, z), \quad (1)$$

where $D_i^p(x, y, z)$ is the dose distribution of the i th beamlet and w_i is its weighting factor. $D_i^p(x, y, z)$ is also referred to as the dose deposition coefficient or the FSPB kernel. According to the methods proposed by Jiang (1998) and Jelen *et al* (2005), a general FSPB dose kernel can be formulated as

$$D^p(x, y, z) = A(\theta, d)F(x', d, z', C), \quad (2)$$

where θ is defined as the angle between the central axis of a beamlet and the central axis of the broad beam, and d is the radiological depth. The patient coordinate system and pencil beam coordinate system used in this work are defined in figures 1(a) and (b), respectively. For the beamlet central axis \overline{SAB} , l is the length of AO' and d is the corresponding radiological depth. $A(\theta, d)$ is a factor that accounts for all off-axis effects. $F(x', d, z', C)$ is the beamlet profile at a depth d with off-axis effects removed. x' and z' are the coordinates in the pencil beam system, defined on the plane perpendicular to the beamlet central axis. C is a set of parameters defining the shape of the beamlet profile (Jiang 1998, Jelen *et al* 2005).

Under the general FSPB framework defined in equation (2), there is flexibility in the selection of a beamlet profile function $F(x', d, z', C)$. The solitary requirement is that this selected function satisfies the self-consistency and normalization conditions. Jiang (1998) derived the FSPB dose kernel as the summation of three error functions. Jelen *et al* (2005) implemented an FSPB model by combining several exponential functions. Lin and his coauthors (Lin *et al* 2006) constructed a simple and finite-term analytic function by using Boltzmann function. The parameters such as C and $A(\theta, d)$ in equation (2) can be dissolved

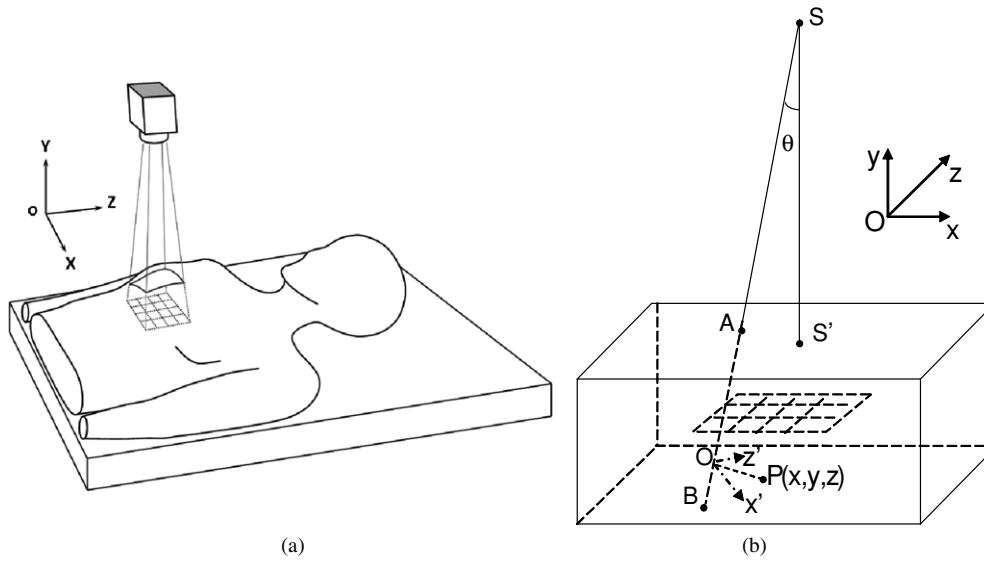


Figure 1. (a) An illustration of the FSPB model and the patient-coordinate system. (b) An illustration of the pencil beam coordinate system for calculating the dose from a beamlet \overline{SAB} to point P . We define the smallest rectangular prism that encases the patient CT images. θ is the angle between the central axes of a broad beam (SS') and a beamlet \overline{SAB} . A is the entrance point and B is the exit point of the beamlet central axis passing through the rectangular prism. O' is the perpendicular projection of point P onto \overline{SAB} . x' and z' are the projection of x and z on the plane perpendicular to \overline{SAB} .

by fitting depth–dose curves and dose profiles at various depths of broad beams. In this study, we use an error function-based FSPB kernel to illustrate our GPU implementation (Jiang 1998):

$$F(x', d, z', C) = \sum_{i=1}^3 \frac{c_1^{(i)}(d)}{4} \left[\operatorname{erf} \left(\frac{c_2^{(i)}(d)a/2 - x'}{\sqrt{2}c_3^{(i)}(d)} \right) + \operatorname{erf} \left(\frac{c_2^{(i)}(d)a/2 + x'}{\sqrt{2}c_3^{(i)}(d)} \right) \right] \cdot \left[\operatorname{erf} \left(\frac{c_2^{(i)}(d)b/2 - z'}{\sqrt{2}c_3^{(i)}(d)} \right) + \operatorname{erf} \left(\frac{c_2^{(i)}(d)b/2 + z'}{\sqrt{2}c_3^{(i)}(d)} \right) \right] \quad (3)$$

where a and b are the side lengths of a beamlet's rectangular cross section; $C = \{c_1^{(i)}, c_2^{(i)}, c_3^{(i)}, \text{ with } i = 1, 2, 3\}$, as a function of d , is a set of parameters obtained by fitting broad beam profiles (Jiang 1998). Note that our implementation allows for this kernel to be easily replaced by any other FSPB kernel.

The dose contribution of a beamlet to a voxel that is far from the beamlet is expected to be very small and can be neglected. To improve computational efficiency, we only compute dose deposition coefficients for voxels inside volumes of interest (VOI). The VOI of a beamlet used in this study is a larger co-central axis beamlet with a cross-sectional side length six times that of the original beamlet (Fox *et al* 2006).

2.2. CUDA implementation

Recently, general-purpose computation on GPU (GPGPU) has been greatly facilitated by the development of graphic card language platforms. One example is the CUDA platform developed by NVIDIA (NVIDIA 2009), which allows the use of an extended C language to

program the GPU. In our work, we use CUDA with NVIDIA GPU cards as our implementation platform.

The GPGPU strategy is well suited for carrying out data-parallel computational tasks, where one program is divided and executed on a number of processor units in parallel. The realization of this single instruction multiple data mode (SIMD) relies on a large number of processing units. General graphic cards such as the GeForce series and the GTX series typically have 32–240 scalar processor units, and the available memory varies from 256 MB to 1 GB. Recently, NVIDIA introduced special graphic computing processors such as the Tesla C1060 that are designed solely for scientific computation. The available memory on the Tesla C1060 is extended to 4 GB. Our implementation is evaluated using a Tesla C1060 card since we believe that it has the optimal trade-off between performance and cost in terms of online re-planning applications.

A GPU has to be used in conjunction with a CPU. The CPU serves as the *host* while the GPU is called the *device*. The CUDA platform extends the concept of C functions to *kernels*. A kernel, invoked from the CPU, can be executed N times in parallel on the GPU by N different CUDA *threads*. For convenience, CUDA threads are grouped to form thread *blocks*, and blocks are grouped to comprise *grids*. The number of blocks and threads has to be explicitly defined when executing a kernel. The threads of a block are grouped into *warps* (32 threads per warp). A warp executes a single instruction at a time across all its threads. On the CUDA platform, the main code runs on the host (CPU), calling kernels that are executed on a physically separate device (GPU). Due to the physical separation of the device and the host, communication between the two cannot be avoided and has to be carefully addressed in CUDA programming.

Figure 2 shows the pseudocode for the implementation of our FSPB model on both CPU and GPU platforms. The CPU implementation (figure 2(a)) describes two loops for dose deposition coefficient calculation. The first is over all beamlets, while the second is over all voxels in VOIs. With this scheme, the computational time for a single beam on a single-core CPU can be estimated as

$$T_{\text{CPU}} = N_{\text{beamlets}} \cdot (T_1 + N'_{\text{voxels}} \cdot T_2), \quad (4)$$

where T_1 is the computational time over a beamlet loop, and T_2 is the time over a loop of a voxel hit by a beamlet. In the GPU pseudocode (figure 2(b)), the two corresponding nested loops from the CPU version are separated into two independent ones (parts I and II) to efficiently use all GPU multiprocessors. The computational time will be reduced to

$$T_{\text{GPU}} = \frac{N_{\text{beamlets}}}{N_{\text{threads}}^{(1)} \cdot N_{\text{blocks}}^{(1)}} \cdot T'_1 + \frac{N_{\text{beamlets}}}{N_{\text{blocks}}^{(2)}} \cdot \frac{N'_{\text{voxels}}}{N_{\text{threads}}^{(2)}} \cdot T'_2, \quad (5)$$

where T'_1 and T'_2 are the computational times when executing on a GPU with a single thread. T'_1 and T'_2 are numerically equal to T_1 and T_2 if the GPU and the CPU have the same clock speed. $N_{\text{threads}}^{(1)}$ and $N_{\text{threads}}^{(2)}$ are the number of threads per block for parts I and II, respectively. $N_{\text{blocks}}^{(1)}$ and $N_{\text{blocks}}^{(2)}$ are the number of blocks for parts I and II, respectively. We need to point out that equation (5) represents the ideal situation. In reality, a GPU has limited numbers of processors and multiprocessors. When the number of threads per block is larger than the warp size (i.e. 32) or the number of blocks is larger than the number of multiprocessors (i.e. 30 for Tesla C1060 GPU), the linear scalability shown in equation (5) will be broken; GPU computation time will not decrease linearly with the increase of the number of threads per block or the number of blocks.

As shown in figure 2(b), besides GPU computation, there are two additional steps, copying data arrays from the host (CPU) to the device (GPU) prior to GPU computation and copying

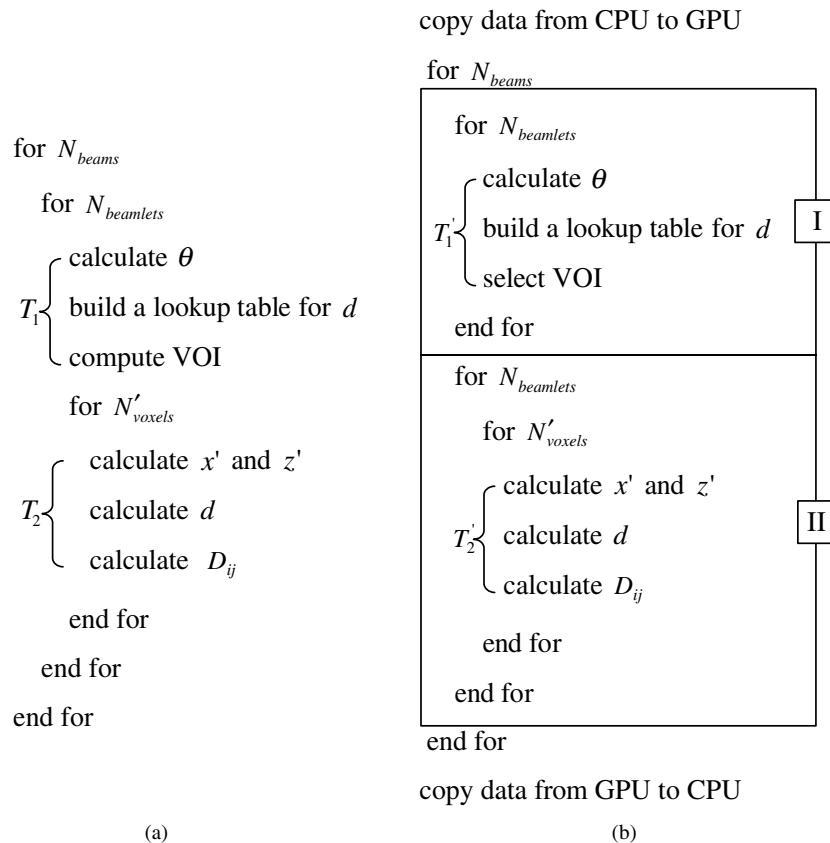


Figure 2. The pseudocode description of our FSPB algorithm implemented on (a) CPU and (b) GPU. Here, the GPU code is parallelized in two different parts. In part I, it is parallelized for all beamlets with both threads and blocks. In part II, the voxels are parallelized with threads and the beamlets are parallelized with blocks. N'_{voxel} refers to the number of voxels contained in a beamlet's VOI.

memory arrays from the device back to the host after computation. Since the data transfer bandwidth within the device ($\sim 73\,500\text{ MB s}^{-1}$ on Tesla C1060) is much higher than that between the device and the host ($\sim 2200\text{ MB s}^{-1}$ on PCI Express device Gen 1), frequent communication between the device and the host can significantly impair the overall efficiency. We thus minimize host–device communication by allocating most of the working variables directly in the GPU memory except those carrying input and output data.

In a typical computation operation, the efficiency of a code is largely determined by the efficiency of memory management. On a GPU, available memory is divided into constant memory, global memory, shared memory and texture memory. The constant memory is cached, which requires only one memory instruction (4 clock cycles) to access. The global memory is not cached, requiring 400–600 clock cycles of memory latency per access. However, the available constant memory is limited to 64 MB on a general GPU card. Due to this limitation, we store only those arrays with constant values, such as the sources' positions, in the constant memory. Optimal usage of global memory requires coalesced memory access, but the radiological depth calculation cannot avoid random data access. To achieve optimized performance, we use the texture fetch feature of CUDA to access data stored in texture memory.

Conceptually, radiological depth is the integral over the passing length of a beamlet central axis, which can be computed using a ray-tracing algorithm such as Siddon's algorithm (Siddon 1985). However, doing ray tracing for each voxel's perpendicular projection O' onto the central axis of each beamlet (figure 1(b)) is highly time consuming. Thus, we first use a trilinear interpolation method to build a radiological depth lookup table. As shown in figure 1(b), the beamlet central axis enters the smallest encasing rectangular prism at point A and exits at point B . We place n equally spaced points on the segment AB and compute the radiological depth from the entrance point A to each of the n points. We can then build a radiological depth lookup table, shown as 'build up a lookup table for d ' in figure 2(a). The radiological depth at perpendicular projection point O' can be obtained with a linear interpolation of two neighboring values stored in the table. Massive memory access occurs when using trilinear interpolation to build the radiological depth lookup table. To avoid this, we place the original CT data in texture memory, which can then be accessed as cached data. The texture fetching possesses a relatively higher bandwidth than global memory access when the coalesced memory access pattern cannot be followed. Another issue worth mentioning here is the hardware implementation of the linear interpolation. CUDA offers linear, bilinear and trilinear interpolation functions at the hardware level. However, due to their insufficient precision, we implemented all interpolation functions in software rather than using the provided hardware functions.

3. Experimental results

3.1. Water phantom

The performance of our FSPB implementation on both GPU and CPU platforms was first evaluated with a water phantom experiment, where five co-planar and equally spaced 6 MV beams of $10 \times 10 \text{ cm}^2$ field size were used to irradiate a $30 \times 30 \times 30 \text{ cm}^3$ water phantom. We used the FSPB model of a 6 MV beam built and validated by Jiang (1998). The performance of our CUDA code was systemically evaluated by independently varying the size of beamlets and voxels. The CPU code was executed on a 2.27 GHz Intel Xeon processor, and the GPU code was run on a NVIDIA Tesla C1060 card. All the testing scenarios are listed in table 1. In each case, the number of voxels involved in computation (i.e. in all VOIs) was recorded. In table 1, T_{CPU} is the sequential execution time with the CPU implementation. T_{GPU}^* and T_{GPU} are the execution times on the parallelized GPU implementation accounting for and excluding CPU–GPU data transfer, respectively. The time for data transfer includes copying data from CPU to GPU before GPU computation and from GPU to CPU after GPU computation. This time can be ignored when estimating the required computational time for online re-planning because the data transfer between D_{ij} calculation, automated target/OAR contouring and plan optimization can occur within the GPU memory. However, when using our GPU implementation for a stand-alone dose calculation, this interval (or at least part of it) should be included in efficiency assessment.

The parallel computation was conducted with CUDA kernels of fixed size. For the parallelization on the beamlet level (figure 2(b), part I), we used 128 threads per block and 256 blocks per grid; for the parallelization on the beamlet and voxel levels (figure 2(b), part II), we used 256 threads per block and the number of beamlets as the block size. Varying the CUDA kernel size had little impact on the execution time T_{GPU} . For example, for case 3 listed in table 1, the computational times were 0.17, 0.18 and 0.18 s when we changed threads number in part II from 64, 128 and 256, respectively.

The computational time on the CPU ranged from 21 to 124 s. On the GPU, excluding data transfer time, the computational time is less than 0.5 s, which leads to a speedup

Table 1. Execution time and speedup factors for different beamlet sizes and voxel sizes. Here, N'_{voxel} is the total number of voxels in all beamlets' VOIs.

No.	Voxel size (cm ³)	Beamlet size (cm ²)	N_{voxel} ($\times 10^6$ voxels)	N_{beamlet} (beamlets)	N'_{voxel} ($\times 10^7$ voxels)	T_{CPU} (s)	T_{GPU} (s)	T_{GPU}^* (s)	Speedup $\frac{T_{\text{CPU}}}{T_{\text{GPU}}}$	Speedup $\frac{T_{\text{CPU}}}{T_{\text{GPU}}^*}$
1	0.50 ³	0.20 ²	0.22	2500	8.64	21.22	0.06	0.11	373.04	192.32
2	0.37 ³	0.20 ²	0.51	2500	1.80	42.80	0.10	0.19	409.34	222.84
3	0.30 ³	0.20 ²	1.00	2500	3.23	78.27	0.18	0.34	419.80	230.22
4	0.25 ³	0.20 ²	1.73	2500	5.27	124.54	0.30	0.56	420.73	223.24
5	0.25 ³	0.25 ²	1.73	1600	4.97	120.14	0.29	0.53	414.85	225.91
6	0.25 ³	0.33 ²	1.73	900	4.70	112.78	0.27	0.46	416.40	244.13
7	0.25 ³	0.50 ²	1.73	400	4.39	100.77	0.24	0.43	417.10	232.61

factor around 400. The computational error for the GPU implementation is defined as $\varepsilon = \frac{1}{N} \sum_{i,j} |D_{ij}^{\text{CPU}} - D_{ij}^{\text{GPU}}|$, where N is the total number of dose deposition coefficients. For all seven cases, the value of ε is constantly around 10^{-6} .

When comparing T_{GPU} to T_{GPU}^* , we observed that the data transfer time between CPU and GPU is comparable to the GPU computational time, which reduces the speedup factor from ~ 400 (T_{GPU}) to ~ 200 (T_{GPU}^*). This observation is not surprising, as it confirms the expensive nature of communication between CPU and GPU. Consequently, our strategy to minimize data transfer by storing data directly on the device appears to be correct.

3.2. A clinical case

The performance evaluation of our FSPB implementation on both GPU and CPU was also carried out for a prostate clinical case, where nine co-planar 6 MV beams were used. The beamlet size was 0.5^2 cm^2 , which yields 528–600 beamlets for each beam. The voxel size was 0.25 cm^3 . For each beamlet, there were about $1.0\text{--}3.0 \times 10^4$ voxels involved in computation. Note that as in the water phantom case, there is essentially no difference between GPU and CPU calculations for this clinical case. The sequential CPU computation took about 4.8 min. For the parallel GPU implementation, it took 0.7 s for the computation only and 1.2 s including data transfer between CPU and GPU. We used the same CUDA kernel size as in the water phantom case.

The accuracy of the radiological depth is related to the distance (or the element length) between two neighboring interpolation points used in the lookup table. To test the accuracy of this lookup table, we sequentially decreased the distance between two neighboring points from the full length of a voxel to one half, one third and one quarter. However, we did not observe any appreciable change in the final values of dose deposition coefficients. This may be due to the relatively small variation in the CT numbers in the prostate case. When applying our implementation to other tumor sites where CT image inhomogeneity is significant, we may need to use a finer element length to build the radiological depth lookup table. For the results presented in this paper, we use half of the voxel dimension as the distance between two neighboring points. We would also like to point out that the GPU time is almost independent of the lookup table resolution; when the element length decreases by a factor of 4, the GPU time only increases by less than 1%.

4. Discussion and conclusions

This paper presents the development of an FSPB-based GPU parallel computing framework for IMRT dose calculation. An analytical FSPB model utilizing three error functions as the FSPB kernel was used for illustration purposes. Any other FSPB models can be easily implemented in our GPU framework by simply replacing a few lines in the CUDA code with an alternative FSPB kernel. FSPB models are particularly suitable for computing dose deposition coefficients for IMRT plan optimization. For most clinical cases, their accuracy is sufficient. However, more advanced dose calculation models may be needed when there exist large inhomogeneities. For this reason, we are also developing a GPU-based parallel platform for Monte Carlo dose calculation.

In this study, the computational speed and numerical accuracy of the GPU code were tested both with a water phantom and a clinical case. The computational accuracy evaluation has shown that the error is well controlled on the level of single floating-point precision of 10^{-6} , which is negligible in real clinical applications.

The achieved speedup gain from using GPUs is affected by a number of factors. First, the speedup depends primarily on hardware configuration. We tested our CUDA code on four different types of NVIDIA GPUs: GeForce 9500 GT, GTX 285, Tesla C1060 and Tesla S1070. Among them, GeForce 9500 GT has only four multiprocessors. Both GTX 285 and C1060 have 30 multiprocessors. S1070 has 4 GPUs, each with 30 multiprocessors. The clock rates of these four cards are very similar varying from 1.3 GHz to 1.48 GHz. When performing the water phantom example case 1 (as shown in table 1) on a single GPU, we found that T_{GPU} values for GTX 285, C1060 and S1070 are nearly identical, which are, however, over seven times faster than that for GeForce 9500 GT. This result confirms that the number of multiprocessors plays a fundamental role in deciding the computation time when the clock rate is fixed.

Another major factor affecting computational speed is the size of GPU memory. GeForce 9500 GT has 512 MB global memory, GTX 285 has 1 GB memory and both C1060 and S1070 have 4 GB memory per GPU. On C1060 and S1070, for all testing scenarios, the computation can be executed in one pass, requiring only two data transfers between the CPU and the GPU, once each at the beginning and end of GPU computation. On GTX 285 and GeForce 9500 GT, for cases 4, 5, 6 and 7 in table 1, the computation cannot be done all at once for five beams. After finishing the calculation of dose deposition coefficients for two or three beams on the GPU, we had to copy the results back to the CPU in order to free the GPU memory for computation over the remaining beams. As mentioned above, the bandwidth between the GPU and the CPU is much lower than that within the GPU, so the additional CPU–GPU data transfer will significantly erode efficiency gains. Choosing the proper GPU for online adaptive therapy is based on a trade-off between cost and performance. Based on the currently available GPUs, we recommend the use of NVIDIA Tesla C1060 cards, which cost around \$1500 and can be readily inserted into a workstation PC.

The speedup factor also depends on the implementation of the FSPB model on CPU. Our CPU implementation is comparable to that of Jelen *et al* (2005). The computation time for the CPU code can certainly be reduced by optimizing the code using various well-established tricks. We did not focus on this issue since for the purpose of this project, what matters the most is the GPU computation time, not the CPU time nor the speedup factor.

The dose deposition coefficient calculation is one of the most significant components in IMRT plan optimization and the main constraint for real-time online adaptive radiotherapy. In this work, the dose deposition coefficient calculation was performed independently on individual beamlets and voxels. The parallel nature of dose calculation warrants significant acceleration by massively parallel computing on GPU. Our GPU implementation of the FSPB model allows the computation of a nine-field prostate IMRT plan within 1 s. The unprecedented speedup factors achieved in this work clearly demonstrates a feasible and promising solution to the realization of real-time online re-planning for adaptive radiotherapy.

Acknowledgments

This work is supported in part by the University of California Lab Fees Research Program. We would like to thank NVIDIA for providing GPU cards and Dr Lai Qi for his help with figures.

References

- Birkner M, Yan D, Alber M, Liang J and Nusslin F 2003 Adapting inverse planning to patient and organ geometrical variation: algorithm and implementation *Med. Phys.* **30** 2822–31

- Bourland J D and Chaney E L 1992 A finite-size pencil beam model for photon dose calculation in three dimensions *Med. Phys.* **19** 1401–12
- de la Zerda A, Armbruster B and Xing L 2007 Formulating adaptive radiation therapy (ART) treatment planning into a closed-loop control framework *Phys. Med. Biol.* **52** 4137–53
- Dubash M 2005 Moore's Law is dead, says Gordon Moore *Techworld* <http://news.techworld.com/operating-systems/3477>
- Fox C, Romeijn H E and Dempsey J F 2006 Fast voxel and polygon ray-tracing algorithms in intensity modulated radiation therapy treatment planning *Med. Phys.* **33** 1364–71
- Fu W H, Yang Y, Yue N J, Heron D E and Huq M S 2009 A cone beam CT-guided online plan modification technique to correct interfractional anatomic changes for prostate cancer IMRT treatment *Phys. Med. Biol.* **54** 1691–703
- Godley A, Ahunbay E, Peng C and Li X A 2009 Automated registration of large deformations for adaptive radiation therapy of prostate cancer *Med. Phys.* **36** 1433–41
- Hissoiny S, Ozell B and Despres P 2009 Fast convolution-superposition dose calculation on graphics hardware *Med. Phys.* **36** 1998–2005
- Jacques R, Taylor R, Wong J and McNutt T 2008 Towards real-time radiation therapy: GPU accelerated superposition/convolution *High-Performance MICCAI Workshop*
- Jaffray D A, Siewerdsen J H, Wong J W and Martinez A A 2002 Flat-panel cone-beam computed tomography for image-guided radiation therapy *Int. J. Radiat. Oncol. Biol. Phys.* **53** 1337–49
- Jelen U, Sohn M and Alber M 2005 A finite size pencil beam for IMRT dose optimization *Phys. Med. Biol.* **50** 1747–66
- Jiang S B 1998 Development of a Compensator Based Intensity-Modulated Radiation Therapy System *PhD Thesis* Medical College of Ohio, Toledo, OH
- Lin H, Wu Y C and Chen Y X 2006 'A finite size pencil beam for IMRT dose optimization'—a simpler analytical function for the finite size pencil beam kernel *Phys. Med. Biol.* **51** L13–5
- Lu W G, Chen M, Chen Q, Ruchala K and Olivera G 2008 Adaptive fractionation therapy: I. Basic concept and strategy *Phys. Med. Biol.* **53** 5495–511
- Meihua L, Haiquan Y, Koizumi K and Kudo H 2007 Fast cone-beam CT reconstruction using CUDA architecture *Med. Imag. Technol.* 243–50
- Mohan R, Zhang X D, Wang H, Kang Y X, Wang X C, Liu H, Ang K, Kuban D and Dong L 2005 Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes *Int. J. Radiat. Oncol. Biol. Phys.* **61** 1258–66
- Noe K O, De Senneville B D, Elstrom U V, Tanderup K and Sorensen T S 2008 Acceleration and validation of optical flow based deformable registration for image-guided radiotherapy *Acta Oncol.* **47** 1286–93
- NVIDIA 2009 *NVIDIA CUDA Compute Unified Device Architecture, Programming Guide* version 2.2 (NVIDIA)
- Ostapiak O Z, Zhu Y and VanDyk J 1997 Refinements of the finite-size pencil beam model of three-dimensional photon dose calculation *Med. Phys.* **24** 743–50
- Preis T, Virnau P, Paul W and Schneider J J 2009 GPU accelerated Monte Carlo simulation of the 2D and 3D Ising model *J. Comput. Phys.* **228** 4468–77
- Riabkov D, Brown T, Cheryauka A and Tokhtuev A 2008 Hardware accelerated C-arm CT and fluoroscopy: a pilot study *Proc. SPIE—Int. Soc. Opt. Eng.* **69132** V-1–9
- Samant S S, Xia J Y, Muyan-Ozcelilk P and Owens J D 2008 High performance computing for deformable image registration: towards a new paradigm in adaptive radiotherapy *Med. Phys.* **35** 3546–53
- Sharp G C, Kandasamy N, Singh H and Folkert M 2007 GPU-based streaming architectures for fast cone-beam CT image reconstruction and demons deformable registration *Phys. Med. Biol.* **52** 5771–83
- Siddon R L 1985 Fast calculation of the exact radiological path for a 3-dimensional CT array *Med. Phys.* **12** 252–5
- Wu C, Jeraj R, Lu W G and Mackie T R 2004 Fast treatment plan modification with an over-relaxed Cimmino algorithm *Med. Phys.* **31** 191–200
- Wu C, Jeraj R, Olivera G H and Mackie T R 2002 Re-optimization in adaptive radiotherapy *Phys. Med. Biol.* **47** 3181–95
- Wu Q J, Thongphiew D, Wang Z, Mathayomchan B, Chankong V, Yoo S, Lee W R and Yin F F 2008 On-line re-optimization of prostate IMRT plans for adaptive radiation therapy *Phys. Med. Biol.* **53** 673–91
- Xu F and Mueller K 2007 Real-time 3D computed tomographic reconstruction using commodity graphics hardware *Phys. Med. Biol.* **52** 3405–19
- Yan D, Lockman D, Martinez A, Wong J, Brabbins D, Vicini F, Liang J and Kestin L 2005 Computed tomography guided management of interfractional patient variation *Semin. Radiat. Oncol.* **15** 168–79
- Yan D, Vicini F, Wong J and Martinez A 1997 Adaptive radiation therapy *Phys. Med. Biol.* **42** 123–32