A Load-Balancing Workload Distribution Scheme for Three-Body Interaction Computation on Graphics Processing Units (GPU)

Ashraf Yaseen¹, Hao Ji², and Yaohang Li^{2*}

¹Department of Electrical Engineering and Computer Science
Texas A&M University-Kingsville, Kingsville, TX 78363
ashraf.yaseen@tamuk.edu

²Department of Computer Science
Old Dominion University, Norfolk, VA 23529
{hji, yaohang}@cs.odu.edu

Abstract

Three-body effects play an important role for obtaining quantitatively high accuracy in a variety of molecular simulation applications. However, evaluation of three-body potentials is computationally costly, generally of $O(N^3)$ where N is the number of particles in a system. In this paper, we present a load-balancing workload distribution scheme for calculating three-body interactions by taking advantage of the Graphics Processing Units (GPU) architectures. Perfect load-balancing is achieved if N is not divisible by 3 and nearly perfect load-balancing is obtained if N is divisible by 3. The workload distribution scheme is particularly suitable for the GPU's Single Instruction Multiple Threads (SIMT) architecture, where particle's data accessed by threads can be coalesced into efficient memory transactions. We use two potential energy functions with three-body terms, the Axilrod-Teller potential and the Context-based Secondary Structure Potential, as examples to demonstrate the effectiveness of our workload distribution scheme.

1. Introduction

Although many molecular simulations are typically confined to evaluating interactions between molecular pairs, recent studies show that three-body or even higher order effects play an important role for quantitatively accurate computation in a variety of molecular simulation applications [1-4]. For example, the three-body effects strongly influence solid-liquid and vapor-liquid equilibrium of fluids [5-8]. The context-based secondary structure potential (CSSP) taking three-body statistical terms into account leads to significant accuracy enhancement in evaluating protein secondary structures [9]. A three-body potential incorporating interaction between a DNA base and a protein residue with regard to the effect of a neighboring DNA base outperforms two-body potentials in specific protein-DNA site recognition [10]. A four-body residue-residue contact potential has demonstrated its effectiveness compared with pairwise potentials in

---

*Corresponding Author.

discriminating native protein conformations[11]. Inclusion of three-body effects in the additive CHARMM protein CMAP potential also results in enhanced cooperativity of α -helix and β -hairpin formations[12].

Despite the advantages of evaluating three-body interactions in molecular simulations, the main obstacle of a potential energy involving three-body or higher order terms is its high computational cost. In general, when external influences are not presented, the potential energy of a system with N particles can be evaluated as

$$E = \sum_{i \neq j} U(p_i, p_j) + \sum_{i \neq j \neq k} U(p_i, p_j, p_k) + \sum_{i \neq j \neq k \neq l} U(p_i, p_j, p_k, p_l) + \dots,$$

where $U(p_i, p_j)$ is the two-body term involving two particles p_i and p_j , $U(p_i, p_j, p_k)$ is the three-body term, $U(p_i, p_j, p_k, p_l)$ is the four-body term, and so on. Three-body potentials explicitly calculate the three-body terms which are ignored in two-body potentials together with other higher order terms. Therefore, provided that the three-body terms are calculated accurately, a three-body potential usually yields better accuracy than the two-body ones. However, the tradeoff is significantly more computing time. The computing time increases largely when higher order terms are included. Generally, two-body terms require $O(N^2)$ operations, while $O(N^3)$ for three-body terms, $O(N^4)$ for four-body terms, and so forth. Computation reduction approaches such as Barnes-Hut method, fast multipole, and particle-mesh[13-16] can significantly reduce the overall computation complexity by simplifying interactions between far apart particles. Nevertheless, simulation using computation involving three- or higher-body terms is still unrealistic for a system with relatively large number of interacting molecules until recent performance improvements have been achieved in computer systems.

Graphical Processing Units (GPUs) have been a powerful computational tool in many science and engineering related problems. Today's GPUs greatly outpace CPUs in arithmetic throughput and memory bandwidth, forming a dramatic shift in the field of high performance computing (HPC) [17]. With the massively parallel computing mechanisms, GPUs are able to deliver performance speedups magnitudes of times higher than the CPU, to solve problems in a few minutes instead of hours or days[18-20]. Hence, GPUs have become the ideal processor to accelerate a wide variety of data parallel applications. Being able to simultaneously calculate forces on N particles, GPU has been employed to accelerate the N -body molecular simulation in a variety of applications. For example, Belleman et al. [21] developed a direct gravitational N -body simulation on GPU. Nyland et al. [22] implemented a fast N -body simulation on GPU in astrophysical applications. Stock and Gharakhani[23] introduced a GPU-accelerated multipole-accelerated treecode method for turbulent flow computations. Lashuk et al. [24] proposed a parallel fast multipole method (FMM) on heterogeneous architectures with CPUs and GPUs. Anderson et al. [25] carried out general molecular dynamics on GPU, by simulating N particles in a finite box with periodic boundary conditions. Friedrich et al. [26] also presented an accelerated molecular dynamic simulation implemented on GPU. Hamada et al. [27] came up with an efficient GPU implementation of Barnes-Hut algorithm on large data sets in astrophysics and turbulence and Jetley et al. [28] provided a scalable implementation on GPU clusters. Grimm and Stadel [29] designed a hybrid symplectic N -body integrator to analysis planet and planetesimal dynamics in the late stage of planet formation with GPU acceleration. Significant speedups have been observed in the above applications.

Although the N -body interactions can be carried out in a straightforward way on a serial processor, efficient parallel implementation to fully take advantage of GPU architectures requires deliberate considerations. Our previous work [20] on symmetric two-body interaction computation

on GPU has led to a workload distribution scheme, which is designed to assign computations of pair-wise atom interactions to GPU threads to achieve perfect or near-perfect load balancing in the symmetric N -body problem while facilitating coalesce memory access. In this paper, we extend our previous GPU-based two-body load-balancing workload distribution scheme to explicitly calculations of three-body terms. The effectiveness of our approach is demonstrated in the computation of the Axilrod-Teller potential [30], a physics-based three-body potential function, and the Context-based Secondary Structure Potential (CSSP), a knowledge-based potential energy function with three-body terms to evaluate protein conformation adopting certain secondary structure pattern[9, 31].

The rest of the paper is organized as follows. Our proposed GPU workload distribution for computing three-body interactions is presented in section 2. Section 3 shows our computational results. Finally, Section 4 summarizes our conclusions and future research directions.

2. GPU-based Load-balancing Scheme for Computing Three-body Interactions

Our load-balancing scheme assumes that the three-body interaction terms are independent of the order of the three particles. In other words, the order permutation of the three particles does not change the potential value. Moreover, for simplicity in illustration, we assume one-thread-per-particle assignment, i.e., each thread keeps one particle information unchanged and then shifts the second or the third particles information at each iteration to obtain all combinations of triplets. A fine-grained assignment other than the one-thread-per-particle assignment to enhance the GPU performance on systems with small number of particles is discussed in the subsection 2.4.

2.1 Serial Implementation

The general serial implementation of three-body interaction computation is straightforward. All one needs to do is to enumerate all triplet combinations of three different particles and then calculate the three-body energy of the triplet. The corresponding pseudo code is illustrated in Algorithm 1. For a molecular system with N particles and assuming that each pair of atoms are interacting, there are totally $N*(N-1)*(N-2)/6$ triplet computations and each particle is involved in exactly $(N-1)*(N-2)/2$ interaction computations. However, for each particle in the outer loop in Algorithm 1, its number of three-body interaction calculations in the middle and inner loops varies gradually from $(N-1)*(N-2)/2$ to 1. Consequently, directly mapping the calculations in the middle and inner loops to GPU threads will lead to a highly unbalanced workload distribution.

Algorithm 1. Serial three-body potential calculation for N particles

```

sumE  $\leftarrow$  0.0; // initialization
fori  $\leftarrow$  1 to N-2 { // first particle, outer loop
    particle1 = i;
    for j  $\leftarrow$  i+1 to N-1 { // second particle, middle loop
        particle2 = j;
        for k  $\leftarrow$  j+1 to N { // third particle, inner loop
            particle3 = k; // sum overall potential
            sumE  $\leftarrow$  sumE + TripleE(particle1, particle2, particle3);
        }
    }
}

```

2.2 Rotational Symmetry

Our implementation of load-balancing workload distribution scheme for three-body interaction computation on GPU is based on the concept of rotational symmetry. Assuming that the N particles in the system are stored in a cyclic array, we use i, j , and k to denote the indices of the three particles in clockwise order, and d_{ij} , d_{jk} , and d_{ki} to denote the position separations between particles i and j , j and k , and k and i , respectively. Here d_{ij} is calculated as

$$d_{ij} = \begin{cases} j - i, & i < j \\ j - i + N, & i > j \end{cases}$$

d_{jk} and d_{ki} are calculated in a similar way. Clearly, we can have the following two properties

- 1). $i \neq j \neq k$, and
- 2). $d_{ij} + d_{jk} + d_{ki} = N$.

Then, we study the position separation pattern of d_{ij} , d_{jk} , and d_{ki} of a triplet(P_i, P_j, P_k). Considering two

position separation patterns d_{ij} , d_{jk} , and d_{ki} and d_{ij}' , d_{jk}' , and d_{ki}' , they are rotationally symmetric if $d_{ij} = d_{ij}'$ and $d_{jk} = d_{jk}'$ and $d_{ki} = d_{ki}'$ or $d_{ij} = d_{jk}'$ and $d_{jk} = d_{ki}'$ and $d_{ki} = d_{ij}'$ or $d_{ij} =$

d_{ki}' and $d_{jk} = d_{ij}'$ and $d_{ki} = d_{jk}'$. Or equivalently, if d_{ij} , d_{jk} , and d_{ki} can turn into d_{ij}' , d_{jk}' , and d_{ki}'

via cyclic rotations, then d_{ij} , d_{jk} , and d_{ki} and d_{ij}' , d_{jk}' , and d_{ki}' are rotationally symmetric. Given an example

of a system where $N = 6$, $\begin{smallmatrix} 1 \\ 3 \text{---} 2 \end{smallmatrix}$, $\begin{smallmatrix} 3 \\ 2 \text{---} 1 \end{smallmatrix}$, and $\begin{smallmatrix} 2 \\ 1 \text{---} 3 \end{smallmatrix}$ are rotationally symmetric but

$\begin{smallmatrix} 1 \\ 3 \text{---} 2 \end{smallmatrix}$ and $\begin{smallmatrix} 1 \\ 2 \text{---} 3 \end{smallmatrix}$ are not. In our GPU implementation, assuming one-thread-per-particle and all threads share the same computation pattern due to GPU's SIMT architecture, rotationally

symmetric position separation patterns indicate that some threads will calculate certain triplets in overlap, which will lead to waste of computational power and, more seriously, erroneous results if not handled correctly. Figure 1 illustrates an example where $N = 6$. Thread 1 starts from particle

P1 and triplet (P1, P2, P3) has position separation pattern $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 1 \end{array}$. If a rotationally symmetric position separation pattern of $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 1 \end{array}$, for instance, $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 1 \text{ --- } 4 \end{array}$ is adopted, then thread 1 will carry out three-body interaction calculation on triplet (P1, P2, P6), which is the same as the three-body interaction calculation on triplet (P6, P1, P2) in thread 6 starting at P6 with position separation pattern $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 1 \end{array}$. In summary, the fundamental idea of our GPU-based algorithm is to uniquely enumerate all position separation patterns where there are no rotationally symmetric pairs.

Figure 1: An example with $N = 6$ with thread 1 starting from P1, and thread 6 from P6. The highlighted particle is the first particle that a thread handles. For separation pattern

$\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 1 \end{array}$, threads 1 and 6 will calculate three-body interactions of triplets (P1, P2, P3)

and (P6, P1, P2), respectively. If thread 1 adopts a position separation pattern $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 1 \text{ --- } 4 \end{array}$,

which is rotationally symmetric to $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 1 \end{array}$, then thread 1 will compute three-body interaction of triplet (P1, P2, P6), which overlaps the computation of thread 6 with

position separation pattern $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 1 \end{array}$.

Algorithm 2. Enumerating all non-rotational symmetric position separation patterns except for the order three symmetric one to calculate three-body interaction of the corresponding triplet particles in a GPU thread.

```

Thread(i, N)                                     // N particles, thread i
{
    localE ← 0.0;                                // initialization
    particle1 ← i;                                // particle 1
    range ← ⌊(N - 1)/3⌋;                          // range of particle 2
    for dij ← 1 to range {
        djk ← dij;
        dki ← N - djk - dij;
        particle2 ← (i + dij) mod N;              // particle 2
        while (dki > dij) {                        // range of particle 3
            particle3 ← (i + dij + djk) mod N;    // particle 3
            localE ← localE + TripleE(particle1, particle2, particle3);
            djk ← djk + 1;
            dki ← dki - 1;
        }
    }
    local_sum_reduce(localE);                      // sum partial scores in local block
}

```

To balance workload distribution among GPU threads, we design a novel workload distribution scheme by enumerating all position separation patterns without sharing rotational symmetry. The workload distribution scheme is described in Algorithm 2. For thread i , the index of the first particle is always i specified by the passed parameter and the second and third particles are selected according to the enumerated position separation patterns. The algorithm enumerates all position separation patterns satisfying

$$d_{ij} \leq d_{jk} \text{ and } d_{ij} < d_{ki}. \quad (1)$$

It is easy to show that any position separation pattern that does not satisfy the above condition is rotationally symmetric to one of the position separation patterns satisfying this condition, because we can always rotate the smallest position separation to d_{ij} . The above condition also indicates that d_{ij} is bounded by $\lfloor (N - 1)/3 \rfloor$. Hence, our algorithm iterates the second particle index from $(i + 1) \bmod N$ to $(i + \lfloor (N - 1)/3 \rfloor) \bmod N$. Then, the third particle is iterated to satisfy (1).

Figure 2 illustrates an example of enumerating triplets by the first thread (thread 1) in a system with 10 particles, using Algorithm 2. For thread 1, the first particle is always P1. The second particle iterates from P2 to P4. When the second particle is P2 ($d_{ij} = 1$), three-body interactions of triplets (P1, P2, P3), (P1, P2, P4), (P1, P2, P5), (P1, P2, P6), (P1, P2, P7), (P1, P2,

P8), (P1, P2, P9) with position separation patterns

$\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 5 \end{array}$, $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 3 \text{ --- } 6 \end{array}$, and $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 2 \text{ --- } 7 \end{array}$ are calculated, respectively. When the second particle is P3 ($d_{ij} = 2$), the three-body interactions of triplets (P1, P3, P5), (P1, P3, P6), (P1, P3, P7), (P1, P3,

P8) with respective separation patterns $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 6 \text{ --- } 2 \end{array}$, $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 5 \text{ --- } 3 \end{array}$, $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 4 \end{array}$, and $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 3 \text{ --- } 5 \end{array}$ are accumulated. When the second particle is iterated to P4 ($d_{ij} = 3$), only one triplet (P1, P4, P7) with

separation pattern $\begin{array}{c} 3 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 3 \end{array}$ can satisfy (1). The completion of the algorithm allows thread 1 to carry out three-body interactions of 12 triplets with different position separation patterns that are not rotationally symmetric. Assuming one-particle-per-thread assignment, the total number of three-body interactions is $12 \times 10 = 120 = 10 \times 9 \times 8 / 6$.

(a): $d_{ij} = 1$, the second particle is P2 and three-body interactions of triplets (P1, P2, P3), (P1, P2, P4), (P1, P2, P5), (P1, P2, P6), (P1, P2, P7), (P1, P2, P8), (P1, P2, P9) with

position separation patterns $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 8 \text{ --- } 1 \end{array}$, $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 7 \text{ --- } 2 \end{array}$, $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 6 \text{ --- } 3 \end{array}$, $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 5 \text{ --- } 4 \end{array}$, $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 5 \end{array}$, $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 3 \text{ --- } 6 \end{array}$, and $\begin{array}{c} 1 \\ \diagup \quad \diagdown \\ 2 \text{ --- } 7 \end{array}$ are calculated, respectively.

(b): $d_{ij} = 2$, the second particle is P3 and three-body interactions of triplets (P1, P3, P5), (P1, P3, P6), (P1, P3, P7), (P1, P3, P8) with separation patterns $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 6 \text{ --- } 2 \end{array}$, $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 5 \text{ --- } 3 \end{array}$, $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 4 \end{array}$, and $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 3 \text{ --- } 5 \end{array}$ are calculated, respectively.

(c): $d_{ij} = 3$, the second particle and only one three-body interaction of triplet (P1, P4, P7) with separation pattern $\begin{array}{c} 3 \\ \diagup \quad \diagdown \\ 4 \text{ --- } 3 \end{array}$ is calculated.

Figure 2: An example of enumerating triplets without sharing rotational symmetry in position separation patterns by the first thread (thread 1) in a system with 10 particles.

The only position separation patterns that Algorithm 2 cannot iterate are order three rotationally symmetric patterns in case of $d_{ij} = d_{jk} = d_{ki}$, when N is divisible by 3. The order three rotationally symmetric patterns will cause the triplets with equal separation distances to be calculated repeatedly by different threads. Figure 3 shows an example of a system with 6

particles, where an order three rotationally symmetric pattern $\begin{array}{c} 2 \\ \diagup \quad \diagdown \\ 2 \end{array}$ exists. As a result, threads 1, 3, 5 over calculate triplet (P1, P3, P5) while threads 2, 4, 6 over calculate triplet (P2, P4, P6). Furthermore, if N is divisible by 3, $(N - 1)*(N - 2)$ is no longer divisible by 6 and thus the total number of $N*(N - 1)*(N - 2)/6$ interaction computations cannot be equally distributed to N threads. Consequently, order three rotationally symmetric patterns require special handling.

Figure 3: Order three rotational symmetry in a system with 6 particles. The highlighted particle is the first particle a thread handles. Threads 1, 3, 5 over calculate triplet (P1, P3, P5) while threads 2, 4, 6 over calculate triplet (P2, P4, P6).

2.3 Load-balancing Workload Distribution Scheme

Algorithm 3 shows the complete workload distribution algorithm with special handling of the case when N is divisible by 3. This algorithm is based on the pseudo-code provided in Algorithm 2. Only the first $N/3$ threads will carry out the three-body interaction computation of triplets with order three rotationally symmetric position separation patterns to avoid over calculation. When N is not divisible by 3, each thread carries out exactly $(N - 1)*(N - 2)/6$ three-

body interaction operations, where perfect load-balancing is achieved. When N is divisible by 3, the first $N/3$ threads carry out an additional iteration of three-body interaction computation for triplets with order three rotationally symmetric patterns. When N is big enough and thereby a lot of iterations are needed, this additional iteration has little impact to the overall system performance and hence we can claim that nearly perfect load-balancing is obtained.

Algorithm 3. Load-balancing workload distribution scheme. Perfect load-balancing is achieved when N is not divisible by 3. When N is divisible by 3, an additional iteration is needed for the first $N/3$ threads.

```

Thread(i, N)                                     // N particles, thread i
{
    localE ← 0.0;                                // initialization
    particle1 ← i;                                // particle 1
    range ← ⌊(N-1)/3⌋;                            // range of particle 2
    for dij ← 1 to range {
        djk ← dij;
        dki ← N - djk - dij;
        particle2 ← (i + dij) mod N;              // particle 2
        while (dki > dij) {                       // range of particle 3
            particle3 ← (i + dij + djk) mod N;    // particle 3
            localE ← localE + TripleE(particle1, particle2, particle3);
            djk ← djk + 1;
            dki ← dki - 1;
        }
    }
    // Special handling 3-way rotationally symmetric of N is divisible by 3
    if (N mod 3 == 0) {                            // N divisible by 3
        if (i <= N/3) {
            dij = djk = N/3;
            particle2 ← (i + dij) mod N;          // particle 2
            particle3 ← (i + dij + djk) mod N;    // particle 3
            localE ← localE + TripleE(particle1, particle2, particle3);
        }
    }
    local_sum_reduce(localE);                      // sum partial scores in local block
}

```

Figures 4(a), 4(b), and 4(c) show the workload distribution when $N = 7, 8$, and 9 , respectively, with perfect load-balancing when $N = 7$ and 8 and nearly perfect load-balancing when $N = 9$. In addition to load balancing, the workload distribution scheme is particularly suitable for the GPU's SIMT architecture [32]. This is due to the fact that, at each iteration step, each thread reads data from different particles with the same stride, which can be coalesced into efficient memory transactions.

	Iteration 1			Iteration 2			Iteration 3			Iteration 4			Iteration 5		
Thread 1	P1	P2	P3	P1	P2	P4	P1	P2	P5	P1	P3	P4	P1	P3	P5
Thread 2	P2	P3	P4	P2	P3	P5	P2	P3	P6	P2	P4	P5	P2	P4	P6
Thread 3	P3	P4	P5	P3	P4	P6	P3	P4	P7	P3	P5	P6	P3	P5	P7
Thread 4	P4	P5	P6	P4	P5	P7	P4	P5	P1	P4	P6	P7	P4	P6	P1
Thread 5	P5	P6	P7	P5	P6	P1	P5	P6	P2	P5	P7	P1	P5	P7	P2
Thread 6	P6	P7	P1	P6	P7	P2	P6	P7	P3	P6	P1	P2	P6	P1	P3
Thread 7	P7	P1	P2	P7	P1	P3	P7	P1	P4	P7	P2	P3	P7	P2	P4

2nd column shift

(a): Perfect Balancing ($N = 7$)

	Iteration 1			Iteration 2			Iteration 3			Iteration 4			Iteration 5			Iteration 6			Iteration 7		
Thread 1	P1	P2	P3	P1	P2	P4	P1	P2	P5	P1	P2	P6	P1	P3	P4	P1	P3	P5	P1	P3	P6
Thread 2	P2	P3	P4	P2	P3	P5	P2	P3	P6	P2	P3	P7	P2	P4	P5	P2	P4	P6	P2	P4	P7
Thread 3	P3	P4	P5	P3	P4	P6	P3	P4	P7	P3	P4	P8	P3	P5	P6	P3	P5	P7	P3	P5	P8
Thread 4	P4	P5	P6	P4	P5	P7	P4	P5	P8	P4	P5	P1	P4	P6	P7	P4	P6	P8	P4	P6	P1
Thread 5	P5	P6	P7	P5	P6	P8	P5	P6	P1	P5	P6	P2	P5	P7	P8	P5	P7	P1	P5	P7	P2
Thread 6	P6	P7	P8	P6	P7	P1	P6	P7	P2	P6	P7	P3	P6	P8	P1	P6	P8	P2	P6	P8	P3
Thread 7	P7	P8	P1	P7	P8	P2	P7	P8	P3	P7	P8	P4	P7	P1	P2	P7	P1	P3	P7	P1	P4
Thread 8	P8	P1	P2	P8	P1	P3	P8	P1	P4	P8	P1	P5	P8	P2	P3	P8	P2	P4	P8	P2	P5

2nd column shift

(b): Perfect Balancing ($N = 8$)

	Iteration 1			Iteration 2			Iteration 3			Iteration 4			Iteration 5			Iteration 6			Iteration 7			Iteration 8			Iteration 9			Iteration 10		
Thread 1	P1	P2	P3	P1	P2	P4	P1	P2	P5	P1	P2	P6	P1	P2	P7	P1	P3	P4	P1	P3	P5	P1	P3	P6	P1	P3	P7	P1	P4	P7
Thread 2	P2	P3	P4	P2	P3	P5	P2	P3	P6	P2	P3	P7	P2	P3	P8	P2	P4	P5	P2	P4	P6	P2	P4	P7	P2	P4	P8	P2	P5	P8
Thread 3	P3	P4	P5	P3	P4	P6	P3	P4	P7	P3	P4	P8	P3	P4	P9	P3	P5	P6	P3	P5	P7	P3	P5	P8	P3	P5	P9	P3	P6	P9
Thread 4	P4	P5	P6	P4	P5	P7	P4	P5	P8	P4	P5	P9	P4	P5	P1	P4	P6	P7	P4	P6	P8	P4	P6	P9	P4	P6	P1	P4	P7	P1
Thread 5	P5	P6	P7	P5	P6	P8	P5	P6	P9	P5	P6	P1	P5	P6	P2	P5	P7	P8	P5	P7	P9	P5	P7	P1	P5	P7	P2	P5	P8	P2
Thread 6	P6	P7	P8	P6	P7	P9	P6	P7	P1	P6	P7	P2	P6	P7	P3	P6	P8	P9	P6	P8	P1	P6	P8	P2	P6	P8	P3	P6	P9	P3
Thread 7	P7	P8	P6	P7	P8	P1	P7	P8	P2	P7	P8	P3	P7	P8	P4	P7	P9	P1	P7	P9	P2	P7	P9	P3	P7	P9	P4	P7	P1	P4
Thread 8	P8	P9	P1	P8	P9	P2	P8	P9	P3	P8	P9	P4	P8	P9	P5	P8	P1	P2	P8	P1	P3	P8	P1	P4	P8	P1	P5	P8	P2	P5
Thread 9	P9	P1	P2	P9	P1	P3	P9	P1	P4	P9	P1	P5	P9	P1	P6	P9	P2	P3	P9	P2	P4	P9	P2	P5	P9	P2	P6	P9	P3	P6

2nd column shift

(c): $N = 9$, due to order three rotational symmetry, only the first three threads handles triplets (1, 4, 7), (2, 5, 8), and (3, 6, 9), respectively, at iteration 10

Figure 4: Workload Distribution Scheme when $N = 7, 8, 9$. Perfect load balancing is achieved when $N = 7$ and 8. Near-perfect load balancing is obtained when $N = 9$.

2.4 Additional Performance Improvement Implementations on GPU

The above load-balancing workload distribution scheme assumes the one-thread-per-particle on GPU architecture. Nevertheless, for molecular systems with small N value, the one-thread-per-particle scheme with N threads may not produce enough threads to fully utilize all resources in GPU. To address this issue, we implement fine-grained threads by dividing the workload originally assigned to one thread to multiple threads so that sufficient threads are produced when N is small. Firstly, we calculate the number of threads that can be assigned to handle interactions computation of a particle (T_p) by dividing the maximum number of threads that a GPU device can launch (T_{max}) over the total number of particles N .

$$T_p = \left\lceil T_{max} / N \right\rceil$$

Then, we distribute the workload of interaction computation to T_p threads. As a result, large number of threads whose total number is near the maximum number of threads that the GPU can launch are created. In addition to fine-grained threads, other standard CUDA programming techniques, including parallel sum reduction, loop unrolling, and coalesce memory access are implemented in order to fully take advantage of the GPU architecture[32].

3. Computational Results

Two three-body potentials, including the Axilrod-Teller potential and the CSSP potential are used to demonstrate the effectiveness of our load-balancing workload distribution scheme on GPU. We name the GPU implementation of Axilrod-Teller and CSSP potential energy functions “GPU-AxT” and “GPU-CSSP,” respectively. The load-balancing workload distribution scheme for three-body interactions is adopted in GPU-AxT and GPU-CSSP. As for the two-body interactions in GPU-CSSP, we used our previous approach to balance the workload [20]. Double precision format is used in both functions. Furthermore, the standard CUDA programming techniques for performance improvement are implemented in both potentials.

The GPU-AxT and GPU-CSSP programs are tested on a server with four Tesla C2070 GPUs. The Tesla C2070 GPU (Fermi architecture) has 14 multiprocessors with 32 cores each, 6GB of global memory, 64 KB of RAM which can be configured between Shared Memory and L1 cache, and 32 KB of registers per multiprocessor.

We benchmark GPU-AxT and GPU-CSSP on a set of systems of various sizes. The GPU time we measured includes the time of transferring the system information (particles) arrays to GPU device memories, GPU execution time, and the time of retrieving the calculated results from GPU. We use the nvcc compiler in CUDA 2.0 with “-O3” flag for GPU implementations.

3.1 Computational Results of Axilrod-Teller Potential

3.1.1 Axilrod-Teller Potential

The Axilrod-Teller potential is an intermolecular potential for the interaction of the van der Waals type between three particles[30]. Considering particles i, j , and k , the Axilrod-Teller potential u_{ijk} is calculated as,

$$u_{ijk} = v \left[\frac{1}{r_{ij}^3 r_{ik}^3 r_{jk}^3} + \frac{3(-r_{ij}^2 + r_{ik}^2 + r_{jk}^2)(r_{ij}^2 - r_{ik}^2 + r_{jk}^2)(r_{ij}^2 + r_{ik}^2 - r_{jk}^2)}{8(r_{ij}^5 r_{ik}^5 r_{jk}^5)} \right]$$

Where α is a non-additive coefficient and r_{ij} , r_{jk} , and r_{ik} are Euclidean distances between particles i and j , j and k , and k and i , respectively.

3.1.2 GPU-AxT performance

We employ GPU-AxT to calculate the Axilrod-Teller potential energy in a particle system where a simulation box of length L is initialized with N Argon particles (atoms). In order to demonstrate the effectiveness of Axilrod-Teller potential implementation on the GPU (GPU-AxT), we run GPU-AxT on simulation boxes of various sizes (L ranges from 14 for 343 particles to 60 for 27,000 particles). The execution time of the Axilrod-Teller potential evaluation is averaged over 100 runs.

Figure 5 shows the performance of GPU-AxT implementations on systems of various sizes in terms of double-precision GFLOPS per second on NVIDIA Tesla C2070. For GPU-AxT implementation with one thread per particle, there are certain inefficiencies for systems with less than 2,500 particles, due to insufficient number of threads to fully take advantage of the GPU architecture. To improve the performance of systems with small number of particles, we adopt a fine-grained implementation by evenly splitting the workload of three-body interaction computations belonging to one thread to multiple threads so that nearly the maximum number of threads that a GPU device can launch is created. Significant performance improvements are found in systems with smaller number of particles when fine-grained threads are employed, whose floating point throughputs are promoted to be closer to those with a lot of particles. As the number of particles increases, the floating point throughput curves of fine-grain threads and one-thread-per-particle start to merge because the increasing number of threads in one-thread-per-particle strategy makes more and more efficient use of the GPU architecture. For systems with more than 3,000 particles, the floating point throughput of GPU-AxT implementation with fine-grained threads is close to but slightly less than that of the two-body potential (262 GFLOPS/s), benchmarked using the “nbody” program provided by Cuda SDK with 100,000 particles. The slight performance loss in three-body potential computation is due to an additional round of memory fetches for triplet computations compared to doublet computations as well as more integer operations to determine the particle indices in each triplet.

Figure 5: Performance comparison of GPU-AxT implementations on systems of various sizes in terms of double-precision floating point throughput on NVIDIA Tesla C2070.

Figure 6 shows the memory-only, arithmetic-only, and full-kernel time that GPU-AxT spends on a system with 3,375 particles. One can find that 70.6% of the memory operations are overlapped with the arithmetic operations. After all, with sufficient number of threads, the memory access latency can be effectively masked.

Figure 6: Memory-only, arithmetic-only, and full-kernel time in GPU-AxT on a system with 3,375 particles. 70.6% of memory operations are overlapped with arithmetic operations.

Our GPU-AxT implementation requires transferring the particles coordinates arrays, from the host memory to the GPU device memory and retrieving the calculated overall energy from the device memory. Table 1 shows the data transfer time ($O_{\text{data-transfer}}$) and the calculation time (t_{kernel}) in small, medium, and large systems. The data transfer time is very small compared to the overall GPU-AxT energy evaluation time.

Table1: Data transferring times and computation time in GPU-AxT

Component	Mode	Time (μs)		
		Small-system (3,375 particles)	Medium-system (15,625 particles)	Large-system (42,875 particles)
Particle coordinate arrays	Host to Device	0.07	0.23	0.62
Axilrod-Teller energy	Device to Host	0.07	0.15	0.32
Axilrod-Teller kernel	On Device	243.35	25,004.82	455,049.73

3.1.3 Load Balancing Scheme vs. Direct Mapping Scheme

Theoretically, assuming every three particles are interacting with each other, if the serial algorithm is directly mapped to GPU implementation, the longest thread needs to carry out $(N - 1) * (N - 2) / 2$ three-body interaction calculations. When the load-balancing scheme is used, each thread handles at most $(N - 1) * (N - 2) / 6 + 1$ three-body interactions. Therefore, the theoretical speedup of the load-balancing scheme over direct mapping scheme is approximately 3. Figure 7 shows that, when no distance cutoff is applied, the GPU-AxT implementation using the load-balancing scheme is around 3 times faster than that of direct mapping. This agrees well with our theoretical analysis. Particularly in practical large-scale simulations that takes hours to days, the three times speedup of load-balancing scheme over direct mapping is much appreciated.

Nevertheless, in practice, when interactions between particles separated over a certain distance are weak enough to be ignored in computation, truncation by distance cutoff is often applied. Provided that the accumulated errors have negligible impact to simulation accuracy and do not affect the stability of the system, using the distance cutoff can significantly reduce the overall three-body interaction computation. Figure 7 also shows that the speedup of the GPU-AxT implementation using the load-balancing scheme when half of box width is used as cutoff distance over that of direct mapping is reduced to approximately 1.8.

Figure 7: Speedup of GPU-AxT Implementation using load-balancing scheme over that of direct mapping. When no cutoff distance is applied, speedup is approximately 3, agreeing well with the theoretical analysis. When half box width cutoff is adopted, speedup is reduced to ~ 1.8 .

The performance reduction of the balanced GPU-AxT with distance cutoff is mainly caused by the divergent branches in the program. Evaluation of the Axilrod-Teller potential with distance cutoff requires testing pairwise particle distances – if any one of the pairwise particle distances is higher than the cutoff distance, the three-body interaction computation will not be carried out. When the threads handling three-body interactions within the distance cutoff as well as those exceeding cutoff co-reside in the same GPU warp, divergent branches will occur in runtime.

Table 2 compares the number of divergent branches in GPU-AxT with half box width cutoff with GPU-AxT without distance cutoff. The performance data is obtained by NVIDIA Compute Visual Profiler 3.2 [33]. When no distance cutoff is adopted, GPU-AxT does not suffer from branch divergence because pairwise particle distances are not necessarily checked against cutoff distance. In contrast, when distance cutoff is applied, branch instructions (if statements) are inserted to compare the pairwise particle distances with the cutoff distance, which potentially leads to divergent branches. When half box width cutoff is used, the divergent branches in GPU-AxT are approximately 15% of the total number of branch instructions, which results in speedup reduction.

Table 2: Comparison of the number of divergent branches in GPU-AxT with half box cutoff distance and GPU-AxT without cutoff.

Box Width	# of Branch Instruction		# of Divergent Branch	
	no Cutoff	w. Cutoff	no Cutoff	w. Cutoff
10	2,628	6,722	0	1,319
12	7,817	16,954	0	1,338
14	39,338	91,987	0	11,766
16	87,558	207,909	0	27,614
18	266,098	640,687	0	89,429
20	500,511	1,182,504	0	152,506
22	1,181,940	2,858,410	0	435,072
24	2,489,780	5,809,260	0	730,359
26	4,829,030	11,573,100	0	1,635,830
28	8,787,680	20,750,600	0	2,717,490
30	15,192,000	36,138,900	0	5,058,260

3.1.4 Applications in Monte Carlo and Molecular Dynamics Simulation

We apply GPU-AxT to a Monte Carlo sampling program and a molecular dynamics program for simulating Argon gas system to demonstrate the effectiveness of GPU-accelerated three-body calculations. The measured computation time is the overall application execution time, which includes the CPU time, the GPU time, and the data transferring time between host memory and device memory.

The Monte Carlo simulation is carried out by using Cartesian all atoms move where the position of every particle in the system is changed by a small random perturbation during a Monte Carlo trial. The Monte Carlo sampling program employs the Metropolis algorithm [34] and the Axilrod-Teller potential energy is evaluated in every iteration step to determine the acceptance of the proposed new conformation. Adopting GPU-AxT, the Monte Carlo sampling program is implemented as a heterogeneous CPU-GPU program. The evaluation of the Axilrod-Teller potential energy is carried out on the GPU while the rest of the Monte Carlo computations are executed on the CPU. We execute the Monte Carlo program using GPU-AxT on a Tesla C2070 simulating a system with 3,375 particles. The computation time is averaged over 1,000 Monte Carlo iteration steps. Compared to that of the direct-mapping implementation (10.94s), the computation time per step in the Monte Carlo program using load-balancing GPU-AxT is reduced to 3.68s – a 2.97 speedup, when no distance cutoff is applied. When half box width distance cutoff is used, computation time per step in the Monte Carlo program using load-balancing GPU-AxT is 0.89s, which is 1.81 times faster than the direct-mapping implementation (1.62s).

The implementation of the molecular dynamics program is based on the pseudocode described in [38], which is tuned and yields 10 times faster than normal implementation according to [38]. The Axilrod-Teller potential is used to estimate the van der Waals interactions between three particles. The main difference compared to the Monte Carlo program, where only the overall potential energy value is needed, is that the forces acting on all particles need to be calculated at each iteration step. In our GPU-AxT implementation for molecular dynamics simulation, we maintain arrays of velocities and accelerations for particles in the system. In each

three-body interaction calculation in GPU-AxT, the accelerations of the three participating particles are calculated as the derivative of their partial potential energy with respect to their positions and are integrated accordingly. Once all interaction calculations are completed in GPU-AxT, the new velocity and position of each particle are then evaluated according to its acceleration.

The molecular dynamic simulation program is also implemented as a heterogeneous CPU-GPU program, where the particle interaction calculations are carried out on the GPU while the rest of the program is executed on the CPU. We carried out the program on a system with 3,375 particles and the computation time is recorded for 1,000 iterations. When no distance cutoff is applied, the overall computation time of the molecular dynamic program using load-balancing GPU-AxT is 11.57 hours, which is about 2.9 times faster than the direct-mapping implementation (33.06 hours). When half box width distance cutoff is used, the overall computation time using load-balancing GPU-AxT becomes 1.91 hours, which is a ~1.5 speedup over the direct-mapping implementation (2.77 hours).

The computational results of the Monte Carlo and molecular dynamics programs are consistent with the analysis in Section 3.1.3. It is important to note that in this work we only consider how the GPU acceleration in three-body interaction calculation using Axilrod-Teller potential can affect the overall performance of the simulation programs. In fact, if more parallel computation components, e.g., the proposal function of generating new particle positions in Monte Carlo or calculations of the new particle velocities and positions in molecular dynamics, are moved to the GPU, more significant performance improvements are possible.

3.2 Computational Results of Context-based Secondary Structure Potential (CSSP)

3.2.1 Context-based Secondary Structure Potential (CSSP)

CSSP is a statistical potential integrating inter-residue interaction potentials for assessing the quality of predicted protein secondary structures[9]. Considering a protein chain with L amino acid residues and a fragment of size S where $S < L$, the CSSP potential of a protein molecule is calculated as,

$$U_{protein} = \sum_{i=0}^{L-S+1} \left(\sum_i^S U(R_i) + \sum_{i \neq j}^S U(R_i, R_j) + \sum_{i \neq j \neq k}^S U(R_i, R_j, R_k) \right) - \sum_{i=1}^{L-S-1} \left(\sum_i^{S-1} U(R_i) + \sum_{i \neq j}^{S-1} U(R_i, R_j) + \sum_{i \neq j \neq k}^{S-1} U(R_i, R_j, R_k) \right),$$

where R_i denotes residue i in the protein chain and $U(R_i)$, $U(R_i, R_j)$, and $U(R_i, R_j, R_k)$ are singlet, doublet, and triplet potential terms, respectively.

3.2.2 Performance of Load-balancing Scheme

Unlike the Axilrod-Teller potential, which is a pure three-body potential, the CSSP potential includes three-body terms together with two- and single-body terms. In GPU-CSSP implementation, the single-body terms are calculated using direct mapping and the two-body terms are calculated using the pairwise load-balancing scheme described in[20], which has a theoretical speedup of ~2.0 over direct mapping scheme. The three-body terms are calculated using the load-balancing workload distribution scheme described in this paper with theoretical speedup of ~3.0 over direct mapping. Figure 8 shows the speedup of the load-balancing GPU-

CSSP over that of direct-mapping on a set of proteins ranged from tens to hundreds of residues, where an average speedup of 2.6 is obtained. This speedup is consistent for small and large proteins.

Figure 8: Performance of the balanced GPU-CSSP with respect to the unbalanced GPU-CSSP on Tesla C2070.

3.3 Scalability of the Load Balancing Scheme across Multiple GPUs

The load-balancing workload scheme can be smoothly scaled across multiple GPUs. Based on the fact that the three-body interaction computations of triplets in each thread are independent, we divide the computation work of each thread into approximately even n_{GPU} partitions, where n_{GPU} is the number of GPUs available, so that each GPU can carry out one partition simultaneously. Once the computations on multiple GPUs are completed, partial sums of the results from each GPU are collected and accumulated. For large-scale three-body computations, the GPU-AxT implementation is compute bound, as memory transaction cost is negligible compared to that of the GPU computations for three-body interaction.

Figure 9 shows that the elapsed time of executing GPU-AxT on a system with 8,000 particles using multiple Tesla C2070 GPUs. One can find that the GPU-AxT implementation can achieve nearly linear speedup by using multiple GPUs over the case with a single GPU.

Figure9: Computational time and speedup of GPU-AxTimplementation using multiple Tesla C2070 GPUs on a molecular system with 8,000 particles. The elapsed time is the average iteration time over 100 runs.

4. Conclusions

In this paper, we investigate the approaches of using GPU to accelerate calculation of general three-body potential energy functions for molecular simulation applications. A workload distribution scheme is developed to achieve perfect load-balancing when N is indivisible by 3 and nearly perfect load-balancing otherwise. The load-balancing workload distribution scheme has demonstrated its effectiveness in the GPU-implementations of theAxilrod-Teller potential and the CSSP potential where three-body terms are incorporated.

The new features of latest GPU architectures have the potential to further improve the performance of ourthree-body energy potential implementations on GPU. For example, the new shuffle (SHFL) instruction of the Kepler GPU architecture allows parallel threads to access data from other threads directly, which enable triplets to be shuffled among threads in a warp and thus enhance parallel efficiency. This will be one of the research directions in our future work.

When a distance cutoff is used, the emergence of potential divergent branches hurts the efficiency of the load-balancing scheme for three-body interaction computation, as shown in Section 3.2.3. Data structures such as quadtree in 2D [35] or octree in 3D [36], may be used to reduce the number of divergent branches and thus enhance computational efficiency, due to the fact that three-body interactions are only needed to be computed in the pre-computed space subdivision. With the cost of auxiliary memory [37], the three-body interaction computation may be accelerated significantly. We will also pursue this direction in our future implementations.

Unfortunately, extending this load-balancing workload distribution scheme to four-body potential calculation on GPUs is not straightforward. The position separation patterns in four-body term calculation are more complicated than those ofthree-bodydue to order two rotational symmetry in position separation patterns when four particles are involved. Order two rotational symmetry does not exist in three-body term calculation and requires deliberate analysis. Nonetheless, in general, the computational cost of four-body or higher order terms are much higher than that of three-body ones while the effectiveness of four-body or higher order terms still need theoretical justifications. Therefore, at this moment, there is no immediate urgency in

extending the load-balancing workload distribution scheme for general four-body or higher order terms calculations on GPUs.

Acknowledgements

We would like to thank the reviewers for their suggestions that greatly help us improve the contents of the manuscript. Yaohang Li acknowledges support from NSF grant 1066471. Hao Ji acknowledges support from ODU Modeling and Simulation Fellowship.

References

1. Elrod MJ, Saykally RJ: **Many-body effects in intermolecular forces**. *Chemical reviews* 1994, **94**(7):1975-1997.
2. Wang L, Sadus RJ: **Influence of two-body and three-body interatomic forces on gas, liquid, and solid phases**. *Physical review E, Statistical, nonlinear, and soft matter physics* 2006, **74**(2 Pt 1):021202.
3. Marcelli G, Sadus RJ: **Molecular simulation of the phase behavior of noble gases using accurate two-body and three-body intermolecular potentials**. *Journal of Chemical Physics* 1999, **111**(4):1533-1540.
4. Marcelli G, Todd BD, Sadus RJ: **Beyond traditional effective intermolecular potentials and pairwise interactions in molecular simulation**. *Computational Science-Iccs 2002, Pt Iii, Proceedings* 2002, **2331**:932-941.
5. Wang L, Sadus RJ: **Three-body interactions and solid-liquid phase equilibria: application of a molecular dynamics algorithm**. *Physical review E, Statistical, nonlinear, and soft matter physics* 2006, **74**(3 Pt 1):031203.
6. Wang L, Sadus RJ: **Effect of three-body interactions on the vapor-liquid phase equilibria of binary fluid mixtures**. *J Chem Phys* 2006, **125**(7):074503.
7. Marcelli G, Sadus RJ: **A link between the two-body and three-body interaction energies of fluids from molecular simulation**. *Journal of Chemical Physics* 2000, **112**(14):6382-6385.
8. Anta JA, Lomba E, Lombardero M: **Exploring the influence of three-body classical dispersion forces on phase equilibria of simple fluids: An integral-equation approach**. *Physical review E, Statistical physics, plasmas, fluids, and related interdisciplinary topics* 1994, **49**(1):402-409.
9. Li Y, Liu H, Rata I, Jakobsson E: **Building a Knowledge-Based Statistical Potential by Capturing High-Order Inter-residue Interactions and its Applications in Protein Secondary Structure Assessment**. *Journal of chemical information and modeling* 2013, **53**(2):500-508.
10. Zhao G, Carson MB, Lu H: **Prediction of specific protein-DNA recognition by knowledge-based two-body and three-body interaction potentials**. *Conf Proc IEEE Eng Med Biol Soc* 2007, **2007**:5017-5020.
11. Feng Y, Kloczkowski A, Jernigan RL: **Four-body contact potentials derived from two protein datasets to discriminate native structures from decoys**. *Proteins* 2007, **68**(1):57-66.
12. Best RB, Mittal J, Feig M, MacKerell AD, Jr.: **Inclusion of many-body effects in the additive CHARMM protein CMAP potential results in enhanced cooperativity of alpha-helix and beta-hairpin formation**. *Biophys J* 2012, **103**(5):1045-1051.
13. Barnes J, Hut P: **A Hierarchical O(N-Log-N) Force-Calculation Algorithm**. *Nature* 1986, **324**(6096):446-449.

14. Darden T, York D, Pedersen L: **Particle Mesh Ewald - an N.Log(N) Method for Ewald Sums in Large Systems.** *Journal of Chemical Physics* 1993, **98**(12):10089-10092.
15. Greengard L: **The rapid evaluation of potential fields in particle systems.** *Thesis (doctoral).* Cambridge, Mass.: MIT Press, Yale University.; 1988.
16. Hockney RW, Eastwood JW: **Computer simulation using particles**, Special student edn. Bristol England ; Philadelphia: A. Hilger; 1988.
17. NVIDIA [<http://www.nvidia.com/page/home.html>]
18. Li Y, Zhu W: **GPU-Accelerated Multi-scoring Functions Protein Loop Structure Modeling.** In: *9th IEEE International Workshop on High Performance Computational Biology.* 2010.
19. Zhu W, Yaseen A, Li Y: **DECMC-GPU: An Efficient Multi-Objective Optimization Method with GPU Acceleration on the Fermi Architecture.** *New Generat Comput* 2011, **29**(2):163-184.
20. Yaseen A, Li Y: **Accelerating knowledge-based energy evaluation in protein structure modeling with Graphics Processing Units.** *J Parallel Distr Com* 2012, **72**(2):297-307.
21. Belleman RG, Bedorf J, Zwart SFP: **High performance direct gravitational N-body simulations on graphics processing units II: An implementation in CUDA.** *New Astron* 2008, **13**(2):103-112.
22. Lars Nyland MH, Jan Prins: **Fast N-Body Simulation with CUDA.** In: *GPU Gems 3.* NVIDIA; 2007.
23. M. J. Stock AG: **Toward efficient GPU-accelerated N-body simulations.** In: *46th AIAA Aerospace Sciences Meeting and Exhibit.* 2008.
24. Lashuk I, Chandramowlishwaran A, Langston H, Nguyen TA, Sampath R, Shringarpure A, Vuduc R, Ying LX, Zorin D, Biro G: **A Massively Parallel Adaptive Fast Multipole Method on Heterogeneous Architectures.** *Commun Acm* 2012, **55**(5):101-109.
25. Anderson JA, Lorenz CD, Travesset A: **General purpose molecular dynamics simulations fully implemented on graphics processing units.** *J Comput Phys* 2008, **227**(10):5342-5359.
26. Friedrichs MS, Eastman P, Vaidyanathan V, Houston M, Legrand S, Beberg AL, Ensign DL, Bruns CM, Pande VS: **Accelerating Molecular Dynamic Simulation on Graphics Processing Units.** *J Comput Chem* 2009, **30**(6):864-872.
27. Hamada T, Yokota R, Nitadori K, Narumi T, Yasuoka K, Taiji M: **42 TFlops Hierarchical N-body Simulations on GPUs with Applications in both Astrophysics and Turbulence.** *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis* 2009.
28. Jetley P, Wesolowski L, Gioachin F, Kale, x, V. L, Quinn TR: **Scaling Hierarchical N-body Simulations on GPU Clusters.** In: *High Performance Computing, Networking, Storage and Analysis (SC), 2010 International Conference for: 13-19 Nov. 2010* 2010. 1-11.
29. Grimm SL, Stadel JG: **The Genga Code: Gravitational Encounters in N-Body Simulations with Gpu Acceleration.** *Astrophys J* 2014, **796**(1).
30. Axilrod BM, Teller E: **Interaction of the van der Waals type between three atoms.** *Journal of Chemical Physics* 1943, **11**(6):299-300.
31. Yaseen A, Li Y: **Context-based features enhance protein secondary structure prediction accuracy.** *Journal of chemical information and modeling* 2014, **54**(3):992-1002.
32. NVIDIA: **CUDA programming guide version 3.1.** In.; 2010.
33. **Compute Visual Profiler 3.2** [developer.download.nvidia.com]
34. Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E: **Equation of State Calculations by Fast Computing Machines.** *Journal of Chemical Physics* 1953, **21**(6):1087-1092.
35. Mazumder P: **Planar Decomposition for Quadtree Data Structure.** *Computer Vision, Graphics, and Image Processing* 1987, **38**: 258-274.
36. Bedorf J, Gaburov E, Zwart SP: **A Sparse Octree Gravitational Nbody Code That Runs Entirely on the GPU processor.** *Journal of Computaitonal Physics* 2012, **231**(7): 2825-2839.
37. Gargantini I: **An Effective Way to Represent Quadtrees.** *Communications of ACM* 1982, **25**(12): 905-910.

38. Marcelli G: **The Role of Three-Body Interactions on the Equilibrium and Non-Equilibrium Properties of Fluids from Molecular Simulation.***Swinburne University of Technology Dissertation* 2001.