

Parallel Iterative Linear Solvers on GPU: A Financial Engineering Case

Abstract—In many numerical applications resulting from computational science and engineering problems, the solution of sparse linear systems is the most prohibitively compute intensive task. Consequently, the linear solvers need to be carefully chosen and efficiently implemented in order to harness the available computing resources. Krylov subspace based iterative solvers have been widely used for solving large systems of linear equations. In this paper, we focus on the design of such iterative solvers to take advantage of massive parallelism of general purpose Graphics Processing Units (GPU)s. We will consider Stabilized BiConjugate Gradient (BiCGStab) and Conjugate Gradient Squared (CGS) methods for the solutions of sparse linear systems with unsymmetric coefficient matrices. We discuss data structures and efficient implementation of these solvers on the NVIDIA’s CUDA platform. We evaluate scalability of our implementations in the context of sparse grid combination technique for pricing multidimensional option pricing PDEs.

Index Terms—Sparse linear iterative solvers, GPU, parallel computing, computational finance.

I. INTRODUCTION

Multidimensional partial differential equation (PDE) problems arise in diverse domains such as mathematical physics, life sciences, and financial engineering, etc. The numerical approximation of these PDEs usually demands computationally intensive solutions and regular tensor product grids have exponential complexity due to high spatial dimensionality of the problem. Full grid solution techniques are usually constrained by the *curse of dimensionality*: the exponential growth in the number of discrete unknowns. The sparse grid combination technique can be used to contain this growth of unknowns in time dependent solutions [18], [3]. In a nutshell, the method discretizes the problem on several sparse grids, then solves these sub-problems which have the same spatial dimensionality as the original problem but coarser discretization, and finally properly combines the partial solutions to get the final one. In the context of this work, we use the technique proposed in [3], which is briefly described in Section V.

Computational efficiency of the sparse grid combination technique depends on the efficient solution of the resulting sub-problems. The linear system solvers can account for a large part of the overall computation time. In this paper we investigate through numerical experiments performance of Krylov subspace based iterative solvers, with which time and memory required per iteration do not increase and no restarting is needed as is the case with GMRes solvers[22]. Iterative methods have additional advantage that they do not change the structure of the problem[21]. In this work, we focus on two most popular iterative solvers, BiCGStab and CGS.

For regular or irregular sparse linear systems, the efficient implementation of solvers on parallel architectures becomes harder. Over the past three decades, a number of research efforts have resulted in parallel sparse linear solvers optimized for latest computational architectures. Our goal is to complement these efforts by developing a comprehensive sparse linear solver package for GPUs or alternatively called GPGPUs. Recently, GPUs have been used in various numerically intensive scientific applications due to their superiority over conventional CPUs with respect to achievable computational power and memory bandwidth. For instance, an Intel QuadCore Xeron 5140 CPU has a peak performance of 29 GFlops, whereas NVIDIA GPU Tesla C1060, has a peak performance of 933 GFlops [34]. Hardware vendors have also provided computational scientists with high level programming tools, like Common Unified Device Architecture (CUDA) from NVIDIA and Stream SDK, a precursor of Close-To-Metal from AMD-ATI, for hiding low level or direct access to GPUs, exposing as massively parallel data parallel processors. Compared to GPUs, CPUs are more flexible and can support a wider range of applications at the cost of greatly increased chip complexity. Specifically, programs that require complicated control flows and large data caches to achieve optimal performance are better suited for CPU-based implementations. On the other hand, contemporary GPUs have a significantly larger number of cores and devote a higher percentage of their transistors to floating point operations. Therefore, a GPU provides massive parallelism and delivers better performance than a CPU for certain applications. The sparse grid computation is an example of this kind, as the computation kernels are local and linear. Demonstrating the effectiveness of linear solvers on the CUDA platform will help us to establish the usefulness of this platform for problems like supernova simulation in Astrophysics, earthquake modelling in Geophysics, or derivatives, also called options, pricing in Finance, which is the focus of this work.

One major challenge in computational finance is the pricing of high dimensional options, called basket or index options, with multiple underlying risk factors [15]. In the simplest multidimensional Black-Scholes model, the number of assets determines the the dimensionality of underlying PDEs. On the one hand, the probabilistic formulation of the pricing problem easily translates into a Monte Carlo algorithm. Monte Carlo methods are flexible and therefore widely used for multidimensional pricing but they suffer from several drawbacks such as a relatively slow convergence and difficulty to compute accurate sensitivities of the solution (known as “Greeks” in finance). On

the other hand, the direct solving of the underlying PDE offers fast convergence and easy computation of the sensitivities, but the method is often prohibitively computationally demanding and suffers from the *curse of dimensionality*. We use sparse grid combination technique described in [3] to solve this problem using GPU based linear solvers.

The structure of the remainder of this paper is as follows. In Section II, we describe the NVIDIA GPU architecture. In Section III, we provide a brief review of works that are of interest with regard to this paper. The efficient parallel implementation details, including sparse matrix storage formats and the CUDA based matrix-vector multiplication (SpMV) libraries, are discussed in Section IV. Having discussed the combination technique in the financial engineering context in Section V, we present numerical and performance results in Section VI. Section VII contains a summary of our findings, and future directions for investigation.

II. NVIDIA GPU ARCHITECTURE AND THE CUDA PROGRAMMING MODEL

In this section, we discuss the GPU parallel computing architecture followed by the CUDA programming model which facilitates the developing data parallel applications on NVIDIA GPUs.

A. Processor Architecture

Traditionally designed to excel in visualization tasks like compute intensive rendering, the architecture of a GPU makes it an ideal candidate for massively parallel data processing. In general, unlike CPU, a GPU has more transistors dedicated to data processing than to data caching and flow control. A basic building block of NVIDIA GPUs is a multiprocessor with 8 cores, up to 16384 32-bit registers, 16KB memory shared between 1024 co-resident threads (a multiprocessor executes a block of up to 16 warps, comprising of up to 32 threads, simultaneously). With up to 240 cores (30 multiprocessors) and memory bandwidth up to 102 GBps, the latest generation of GPUs offers extremely cost-effective computational power not only for visualization but also for general purpose scientific computations.

B. Memory Architecture

NVIDIA GPU memory model is highly hierarchical and there exist per-thread local memory, per-thread-block shared memory and device memory which aggregates global, constant and texture memory allocated to a grid, an array of thread blocks. A thread executes a kernel, GPU program, and communicates with threads in the same thread block via high-bandwidth low-latency shared memory. Generally, optimizing the performance of CUDA applications could involve optimizing data access patterns to these various memory spaces. Each of the memory space has certain performance characteristics and constraints. Efficient implementation of solver kernels must consider carefully CUDA memory spaces, specifically, local and global memories which are not cached and have high access latencies.

C. CUDA programming model

NVIDIA's Compute Unified Device Architecture (CUDA) is a general purpose parallel computing architecture with a novel parallel programming model and instruction set architecture. C for CUDA exposes the CUDA programming model as an abstraction of GPU parallel architecture using a minimal set of extensions to the C language by allowing programmers to define C functions, called kernels. When called, these kernels are executed N times in parallel by N different CUDA threads in a hierarchical fashion, as opposed to only once as is the case in regular C functions. NVIDIA also provides CUBLAS, a BLAS (Basic Linear Algebra Subprograms) library ported to CUDA, which enables the use of GPUs without direct operation of the CUDA drivers. Our solver implementations use this library mainly for vector-vector operations. CUBLAS does not provide sparse matrix storage structures.

III. RELATED WORK

GPU based linear solvers: GPU memory can be efficiently utilized for solvers where the matrix has a regular structure. In this work, our target is to solve systems with irregular sparsity. A few other algorithms have been studied to solve sparse and dense linear systems. Dense linear algebra routines are provided by NVIDIA, and their careful optimizations are studied in [29]. It is well known that due to their regular access patterns, dense linear algebra algorithms are well suited to GPU architecture. However, due to the amount of data contained in a dense matrix, in most cases the computations are bandwidth limited. Direct factorization-based solvers have been ported to GPUs [30], [31]. Most of these works rely on blocking strategies to parallelize the operations.

Sparse linear algebra is somewhat more difficult to adapt to GPUs, at least for unstructured problems. Several techniques have been proposed in the literature [27], [13], [25]. The major issues involve how the sparse matrix is stored (compressed storage formats), and whether blocking is used. To exploit massive parallelism offered by GPU, the optimizations for reducing memory footprint and hence hiding the memory access latency is very important. Bell et al [25] compare GPU SpMV results with SpMV results on various multi-core platforms obtained by [26] and illustrate that GPUs offer best performance.

The first GPU-based Conjugate Gradient solver for unstructured matrices is proposed in [27]. To utilize memory bandwidth, blocked compressed sparse row (BCSR) format matrix storage is used in [13] instead of CSR. BCSR decreases number of memory fetches from the device memory to some extent, however, number of elements to be multiplied increases. Both these works solve systems in single precision floating point. A mixed precision, multi-grid solver for a GPU cluster is proposed in [33]. The multi-gpu based general purpose symmetric linear systems solver with double precision solution accuracy is presented in [32].

In [27], authors use textures to store non-zero coefficients of a matrix and its associated two level look tables for CSR format to implement conjugate gradient solver. The lookup

Kernels		Methods		
Kernel	FLOPs	Method	Initial	Iterations
SpMV	2α	CGS	$10n + 2\alpha - 4$	$24n + 4\alpha - 3$
dot	$2n - 1$	BiCGStab	$10n + 2\alpha - 4$	$30n + 4\alpha - 5$
axpy	$2n$	BiCG	$5n + 2\alpha - 1$	$15n + 2\alpha - 4$
scal	n			

TABLE I

FLOPS FOR BASIC KERNELS AND ITERATIVE SOLVERS. n IS DIMENSION OF A MATRIX, α IS NUMBER OF NONZEROS.

table is used to address the data and to sort the rows of the matrix according to the number of non-zero coefficients in each row. Then an iteration is performed on the GPU simultaneously over all rows of the same size to complete, for instance, a matrix-vector operation. Another approach to implement sparse matrices based on CSR format was proposed in which vertex buffer is used to store non-zero elements [28]

Sparse Grid methods for option pricing: Combination technique was first introduced in [3]. Some results on convergence and error analysis of the method (which are not under investigation here) can be found in [1], [6], [7], [10]. The sparse grid combination technique has been successfully applied in fluid mechanics [4], [5], in data mining [8] or in finance [2], [12]. In this paper, we deal with the combination technique, for solving the Black-Scholes PDEs, which is presented in [12]. In this paper authors also discuss an approach where individual sparse grids are solved in parallel on a cluster of workstations.

IV. ITERATIVE SOLVERS ON GPU

We are interested in methods for solving unsymmetric systems of linear equations arising in the sparse grid methods. One of the leading families for linear system solvers is iterative solvers known as Krylov subspace methods [20], [21], [23]. We selected BiCGStab and CGS methods for their suitability for solving unsymmetric linear systems.

Iterative methods use four basic computational kernels: matrix vector products, preconditioner, inner product (dot), and vector update (axpy). The choice of the preconditioner is very important for the efficient solution of a linear system [11], but we will not discuss preconditioning here because it is often problem dependent. The constituent kernels along with their computational costs are presented in Table I. The efficiency of any iterative method is determined primarily by the performance of the matrix-vector products and therefore on the storage scheme used for the matrix. We elect some sparse matrix representations which could be suitable for matrices resulting in the combination technique and evaluate the performance of linear solvers with respect to various implementations of GPU based matrix-vector products.

A. Sparse Matrix Formats

There exist several sparse storage formats with the aim of representing sparse matrices economically. These differ in

terms of amount of storage required, the accessing methods such as the amount of indirect addressing required for fundamental operations like matrix-vector products, and their adaptability for parallel architectures. A survey of various sparse storage formats can be found in [20]. Due to matrix sparsity, memory access patterns tend to be highly irregular and utilization of global uncached memory can suffer from low spatial or temporal locality. Each format takes advantage of specific properties of the sparse matrix and may achieve different degree of efficiency of space and computational efficiency. We prefer to consider general storage formats which are suitable for matrices with arbitrary sparse structure. We consider following formats in our study:

- **Coordinate (COO):** is a general sparse matrix format that comprises of arrays *row*, *col* and *data* to store row indices, column indices and values of nonzero matrix entries, respectively. This format is very space inefficient and computationally intensive among the other formats we considered.
- **Compressed Sparse Row (CSR):** is a general-purpose sparse matrix format. It does not consider any ordering among nonzero values within each row. Subsequent nonzeros of rows are stored in contiguous memory, and additional integer arrays specify column index for each nonzero and beginning of offset of each row.
- **Block Compressed Sparse Row (BCSR):** is particularly useful when the sparse matrix has square dense blocks of nonzeros in some regular pattern. It enables register blocking strategies, and vector processing significantly reduces the required memory bandwidth and computational time for matrices with large block sizes [20].
- **Hybrid HYB:** is a combination of the **Ellpack-Itpack (ELL)** (or **Diagonal (DIA)**) and COO format, by coupling the speed of ELL (or DIA), utilizing the memory bandwidth efficiently, and the flexibility of COO. It is usually the fastest format for a wide range of unstructured matrices.

B. Sparse Matrix Operations

Sparse matrix-vector multiplication is arguably the most important operation in sparse matrix computations. Iterative solvers generally require hundreds, if not thousands, matrix-vector products to reach convergence. Over the past decade, there has been significant amount work on optimizing SpMV. Most of the work has focused on optimizing sparse matrix kernels on general-purpose architectures. SpMV being a memory bound kernel, most optimizations target performance improvements at various memory levels in the memory hierarchy. The optimizations include optimal data structure for storing sparse matrix, exploiting block structures in sparse matrix and blocking for reuse at the level of cache, TLB and registers [14]. Various optimizations have been proposed taking into account the complex memory hierarchy and unconventional mapping of computation to the coresident threads on GPUs.

Let us take an example of a sparse matrix A , as represented in Figure 1(a), its CSR representation in 1(b) and a dense

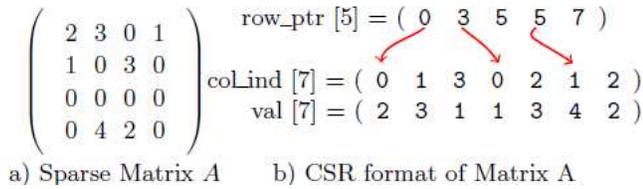


Fig. 1. An example of CSR representation.

vector x of length n . The objective is to compute $y = Ax$, where y is the output dense vector of length n . In Figure 2, we present a sequential implementation of the SpMV procedure using CSR format. There are several ways to parallelize this procedure. A simple translation of this sequential routine into a CUDA kernel is given in Figure 3. In this naive approach of parallelizing the main loop, each row is assigned to one thread. One of the major problems with this parallel implementation on GPU is the way how threads within a warp access the *col_ind* and *val* arrays. Despite these values are stored contiguously, each thread reads the elements of its row sequentially, resulting in non-contiguous accesses to off-chip memory and thus hampering the overall performance of the memory bound kernels. Additionally, if the nonzero elements are unevenly distributed across all the rows, it may lead to poor utilization of the resources keeping many threads idle. Several algorithmic as well as architecture specific optimizations have been proposed to improve the performance of both serial and parallel implementations [26], [25], [14], [13]. We experiment with the following three libraries to implement the linear solvers. We use CUBLAS for dense vector operations.

```

__host__ void
spmv_csr_serial ( const int num_rows ,
                  const int * row_ptr ,
                  const int * col_ind ,
                  const float * val ,
                  const float * x ,
                  float * y )
{
  for ( int row = 0; row < num_rows ; row ++){
    float rowsum = 0;

    int row_start = row_ptr [row];
    int row_end = row_ptr [row +1];

    for (int jj = row_start ; jj < row_end ; jj ++ )
      rowsum += val [jj] * x[col_ind[jj]];

    y[row] += rowsum;
  }
}

```

Fig. 2. Serial routine for CSR SpMV

1) *NVIDIA SpMV Library*: NVIDIA made available a library for SpMV computation and is one of the prominent efforts in this direction. In [25] authors discuss the implementation details of various sparse matrix formats and their efficient representation on NVIDIA GPUs. The library supports DIA and ELL formats suitable for small sparse/dense

```

__global__ void
spmv_csr_scalar_kernel ( const int num_rows ,
                         const int * row_ptr ,
                         const int * col_ind ,
                         const float * val ,
                         const float * x ,
                         float * y )
{
  int row = blockDim.x * blockIdx.x + threadIdx.x ;

  if(row < num_rows ){
    float rowsum = 0;

    int row_start = row_ptr [row];
    int row_end = row_ptr [row +1];

    for (int jj = row_start ; jj < row_end ; jj ++ )
      dot += val[jj] * x[col_ind[jj]];

    y[row] += rowsum;
  }
}

```

Fig. 3. A naive CUDA kernel for CSR SpMV

matrices. The other formats suitable for large sparse matrices include COO, CSR, HYB and Packet (PKT) format.

The *COO-vector* kernel based on segmented reduction is robust with respect to variations in row sizes. It is reliable and complements the deficiencies of the other SpMV kernels. *COO-scalar* kernel uses only one thread and thus exhibits behaviour of a sequential algorithm. The *CSR-scalar* kernel uses one thread per matrix row. It has low bandwidth utilization and hence poor computational performance, as it does not exploit memory coalescing. On the contrary, *CSR-vector* uses a 32-thread warps per row results in contiguous memory access. This may result in poor load balancing by increasing the number of idle threads when the number of nonzero coefficients is less than the warp size. Performance of *CSR-scalar* is rarely competitive with alternative choices, while the vector kernel efficiently handles large row sizes. The *HYB kernel* is a combination of *COO-vector* and ELL kernels. All the kernels that we used in our implementation can also benefit from the read-only *texture cache* present on all CUDA-enabled devices.

2) *IBM SpMV Library*: In [14], authors propose GPU specific, application oriented optimizations for efficient execution of SpMV kernels using CUDA. The GPU specific optimizations include i) Exploitation of synchronization-free parallelism for achieving intra-thread-block synchronization across the rows, ii) Optimized thread mapping based on the *affinity towards optimal memory access pattern*, iii) Enabling hardware optimized global memory coalesced accesses and iv) Exploiting data reuse of input vector elements by caching the elements in on-chip memories like texture (hardware) cache, registers or shared memory (software) cache. These optimizations are for the CSR format and the authors claim to achieve better performance than the NVIDIA SpMV library. We evaluate the performance of our solvers with Padded CSR

format, in which zeros are padded to ensure that the number of entries in each row is a multiple of 16 in order to achieve aligned global memory access. Texture memory can also be used to store the input vector to utilize the read-only texture cache in order to achieve performance gains due to input vector reuse.

3) *CNC SpMV Library*: In [13], authors present SpMV package which includes efficient implementation for BCSR format. The sparse storage format groups non-zero values in blocks of size $BN \times BM$ in order to maximize the memory fetch bandwidth of GPUs, to take advantage of registers to avoid redundant fetches (register blocking), and to reduce the number of indirections due to the reduced size of lookup tables. Although each indirection results in dependent memory fetches, it introduces memory latencies that need to be hidden by the GPU to achieve a good efficiency. We experiment with implementations of 2×2 and 4×4 blocks for lower and higher filling ratio.

V. COMBINATION TECHNIQUE FOR SOLVING THE BLACK-SCHOLES PDE

In this section, we discuss the financial test case chosen for evaluating the GPU implementations of linear solvers. The objective is to compute the price of European derivatives on a d -dimensional basket of risky assets in the multidimensional Black-Scholes framework.

Let us denote $S_i(t)$ the price of the i -th asset, $\sigma_i(t)$ its volatility, $\rho_{ij}(t)$ the linear correlation between assets i and j , and r the risk-free interest rate. We consider a European option with maturity T and payoff $h(\mathbf{S})$. The value $V(\mathbf{S}, t)$, $\mathbf{S} \in \mathbb{R}_+^d$, $t \in [0, T]$ of this option is solution of the following Black-Scholes partial differential equation:

$$\frac{\partial V}{\partial t} = -\frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d \rho_{ij} \sigma_i \sigma_j S_i S_j \frac{\partial^2 V}{\partial S_i \partial S_j} - r \sum_{i=1}^d S_i \frac{\partial V}{\partial S_i} + rV, \quad (1)$$

with terminal condition $V(\mathbf{S}, T) = h(\mathbf{S})$. With constant parameters, this equation can be reduced to a multidimensional heat equation. Although very academic, this test case is similar to the PDEs of more realistic models. See [9] for a large panel of PDE in finance.

As mentioned earlier, solving PDE (1) in a high-dimensional case (say $d > 3$) is limited by memory restriction since the problem is affected by the so-called ‘‘curse of dimensionality’’: on a regular full grid of mesh size 2^{-n} , the discretized problem grows exponentially with the dimension. The sparse grid combination technique allows to solve for higher dimension by reducing the size of the problem. The framework in our case is as follows. For simplicity, we assume our derivative contract is a basket up-and-out barrier option. In this case, the value of the option is zero as soon as one of the assets crosses a given value B , so that the PDE (1) can be solved with homogeneous Dirichlet boundary condition. Let us denote $\Omega = [0, B]^d$ as our domain. For each multi-index $\mathbf{i} \in \mathbb{N}^d$, we denote $\Omega_{\mathbf{i}}$ the Cartesian sparse grid on that domain with mesh

d	MC interval	l=6	7	8	9	10
2	[15.97-16.04]	16.68	16.49	16.10	16.06	15.99
3	[20.87-20.94]	23.49	22.59	21.69	20.80	20.85
4	[24.63-24.69]	26.98	31.43	NA	NA	NA
5	[27.61-27.68]	19.6145	NA	NA	NA	NA

TABLE II
NUMERICAL RESULTS. MONTE CARLO CONFIDENCE INTERVALS COMPUTED USING $1E6$ PATHS AND 25 TIMESTEPS. NOTE THAT MC PRICES ARE POSITIVELY BIASED (SEE TEXT).

size 2^{-i_k} in the k -th dimension. The standard finite differences solution of PDE (1) on the anisotropic grid $\Omega_{\mathbf{i}}$ is denoted $\tilde{V}_{\mathbf{i}}$. Then the solution V_n^{CT} obtained at level n with the sparse grid combination technique is:

$$V_n^{CT} = \sum_{l=n}^{n+d-1} (-1)^{d-1-l+n} \binom{d-1}{l-n} \sum_{i_1+\dots+i_d=l} \tilde{V}_{\mathbf{i}} \quad (2)$$

Therefore, the sparse grid combination technique allows to build a level n discretization of the problem of size $\mathcal{O}(2^n n^{d-1})$ (on level n) instead of size $\mathcal{O}(2^{nd})$ in the full grid case. It also offers a natural coarse parallelization, since all solutions $\tilde{V}_{\mathbf{i}}$ summed in equation (2) can be computed independently from one another. All sub-grids $\Omega_{\mathbf{i}}$ are solved sequentially a *single* GPU using BiCGStab and CGS solvers. Therefore, the speedups we are present in Section VI could be enhanced with a multiple GPU implementation using this ‘‘natural’’ parallelization.

VI. NUMERICAL RESULTS AND PERFORMANCE

In this section, we discuss the results of a set of experiments using NVIDIA GPU to demonstrate the performance behavior of BiCGStab and CGS.

The sequential, double precision versions of BiCGStab and CGS solvers were developed and experimented on Intel Xeon E5420, see Table IV. For the dense vector operations we used BLAS library (1.2-1.3ubuntu3) and supported only CSR sparse format. The sequential solvers and any other CPU bound code was compiled using gcc compiler at $-O3$ optimization level. The GPU based solvers were executed using Tesla C870 (G80 GPU) with configuration presented in Table IV, connected to a host x86/Linux system through 16-x PCI Express bus. The CUDA kernels for SpMV operations were compiled using NVIDIA CUDA compiler 2.1 (nvcc) to generate the device code. The device code was compiled with $-arch=sm_10$ flag and $-O3$ optimization level. Note that due to limitations of Tesla C870, experiments for GPU based solvers were done with single precision arithmetic.

We consider the test case presented in Section V, with an up-and-out best-of barrier option on d assets. The payoff is written $h(\mathbf{S}) = (\max_i S_i - K)^+$ on the domain $[0, B]^d$. Our numerical results are presented in Table II. The following parameters have been used to obtain these results: $\sigma_i = 0.20$, $\rho_{ij} = 0.0$, $r = 0.03$, $T = 1$, $K = 100$, $B = 200$. We solve the PDE using the combination technique at level

d	Level	Total Time(s)	Solver Time(s)
3	6	15.71	15.10
3	7	70.26	65.30
3	8	283.08	242.13
3	9	1155.27	818.10
3	10	5375.96	2553.31

TABLE III

TIME IN SECONDS FOR PRICING A BASKET OF THREE EQUITIES WITH DIFFERENT REFINEMENT LEVELS.

Feature	C870	E5420
Multiprocessors	16	2
Processor cores	128	8
Processor Clock	1.35 GHz	2.6 GHz
Off-chip Memory Size	1.5 GB	4 GB
Peak Performance	500 GFlops	80 GFlops

TABLE IV

ARCHITECTURAL CONFIGURATIONS OF NVIDIA TESLA C870 AND INTEL XEON E5420.

$n = 6, \dots, 10$ with 5000 timesteps. The accuracy of our numerical results is verified by comparison with a Monte Carlo 95% confidence interval obtained with average of 100 prices each computed with $1E6$ simulation paths and 25 timesteps. It is known that these confidence intervals are positively biased, since the evaluation of the crossing of the barrier is not evaluated continuously. Table II shows that convergence is achieved, at least for $d = 2, 3$. For error analysis of the combination technique please refer to [7]. We present the resulting linear systems for $d = 3$ and $n = 10$ in Table V. For lower dimensions and levels the linear systems are smaller in size, however, the number of unknowns increases with the increase in dimensions and approximation level. For instance, in eight dimensions a coarsest reasonable sparse grid can have up to 200 million unknowns [3].

For iterative solvers in the experiments, the convergence tolerance of the iterative methods is set to $1E-6$. The execution times using sequential solvers for $d = 3$ with $n = 6, \dots, 10$ are presented in Table III. The column *Solver Time* represents the time required for solving the linear discrete systems. The remaining of the *Total Time* is spent on constructing partial grids and combining the subsolutions. It usually remains constant and is independent of parallel or sequential solvers used. Hence, for the performance evaluation, we only consider the *Solver time*. The combined speedups achieved for solving all the grids using GPU based BiCGStab and CGS are presented in Figure 4 and Figure 5 respectively. The speedup is computed as the ratio of the parallel execution time (T_p) using Tesla and the execution time using one CPU (T_1). Overall performance of these solvers in terms of mega FLOPs (MFLOPs) is presented in Figure 6 and Figure 7. The total number of FLOPs are computed as described in Table I.

We evaluate the performance of BiCGStab and CGS with respect to the underlying SpMV kernels. For the analysis, we call both the solvers by the SpMV kernel name used (e.g. *CSR-*

$\langle l_1, l_2, l_3 \rangle$	N	NNZ	$\langle l_1, l_2, l_3 \rangle$	N	NNZ
$\langle 8,1,1 \rangle$	255	763	$\langle 5,3,1 \rangle$	217	1009
$\langle 7,2,1 \rangle$	381	1645	$\langle 5,2,2 \rangle$	279	1563
$\langle 6,3,1 \rangle$	441	2065	$\langle 4,4,1 \rangle$	225	1065
$\langle 6,2,2 \rangle$	567	3195	$\langle 4,3,2 \rangle$	315	1863
$\langle 5,4,1 \rangle$	465	2233	$\langle 3,3,3 \rangle$	343	2107
$\langle 5,3,2 \rangle$	651	3895	$\langle 6,1,1 \rangle$	63	187
$\langle 4,4,2 \rangle$	675	4095	$\langle 5,2,1 \rangle$	93	397
$\langle 4,3,3 \rangle$	735	4627	$\langle 4,3,1 \rangle$	105	481
$\langle 7,1,1 \rangle$	127	379	$\langle 4,2,2 \rangle$	135	747
$\langle 6,2,1 \rangle$	189	813	$\langle 3,3,2 \rangle$	147	847

TABLE V

A SET OF DISCRETE SYSTEMS FOR $d = 3$ AND $l = 10$ WITH LEVEL OF DISCRETIZATION IN EACH DIRECTION GIVEN BY THE PERMUTATION OF $\langle l_1, l_2, l_3 \rangle$. N IS DIMENSION OF A MATRIX AND NNZ IS NUMBER OF NONZEROS. THE TOTAL NUMBER OF DISCRETE SYSTEMS IS 85.

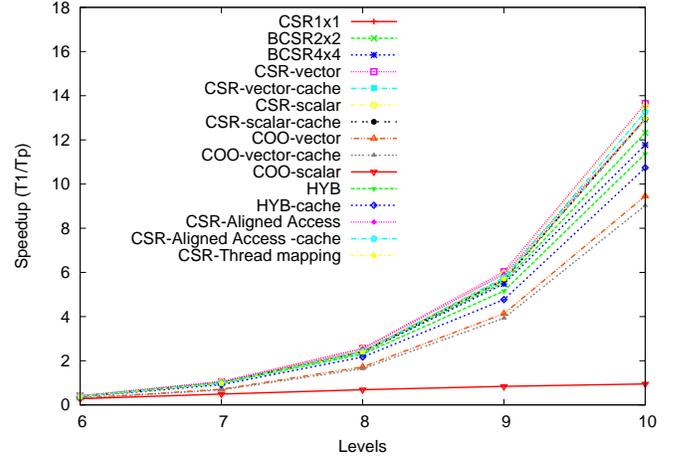


Fig. 4. Speedup for BiCGStab using GPU for increasing problem sizes.

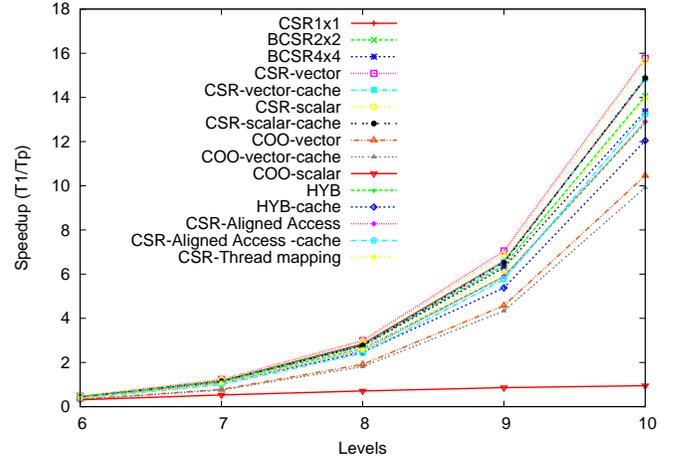


Fig. 5. Speedup for CGS using GPU for increasing problem sizes.

Scalar stands for a BiCGStab or CGS solver using NVIDIA's scalar CSR SpMV kernel). The performance characteristics of both the solvers follow similar behavior for scalability, hence in the following discussion we do not distinguish between both the solvers, if not mentioned otherwise. The other solver kernels (*dot*, *axy* and *scal*) in both the solvers are assumed

to take equivalent time for solving a particular grid.

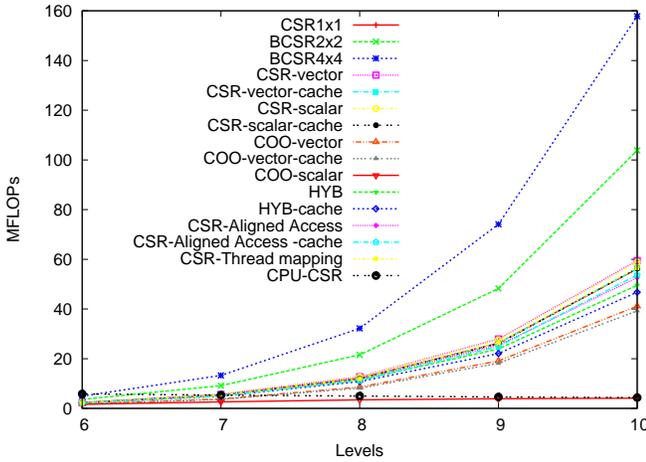


Fig. 6. Overall speed (megaflops) of BiCGStab using GPU.

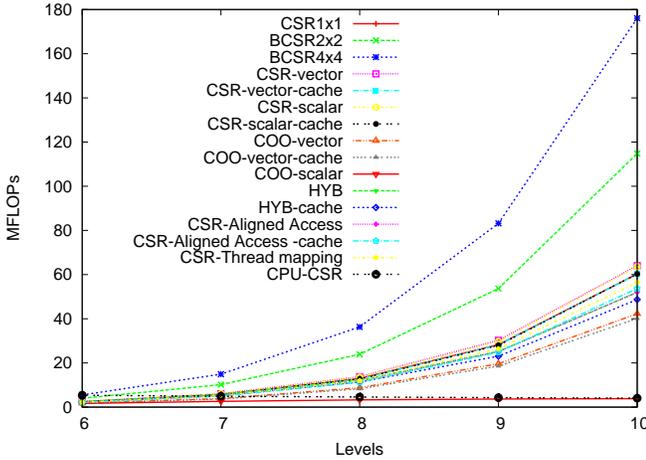


Fig. 7. Overall speed (megaflops) of CGS using GPU.

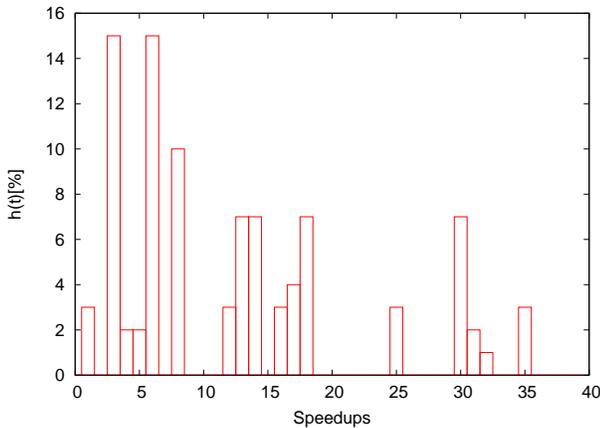


Fig. 8. Histogram $h[t]$ of Speedups achieved by GPU based BiCGStab using *NVIDIA-CSR-vector* kernel for solving discrete systems represented by Table V.

With *CNC library*: *BCSR2x2* and *BCSR4x4* kernels take advantage of square dense block patterns in a sparse matrix, saving in locations and reduction in indirect addressing. The sparse grids resulted in the combination technique does not exhibit any of such pattern. This format leads to register level data reuse resulting coarse grained parallelism, however in our case, suffers from non-optimal global memory accesses. This attributes to its poor performance compared to other variants of CSR formats that we investigated. *CSR1x1* is uses one thread per row and 16 threads per block. The equivalent *NVIDIA CSR-scalar* kernel uses 256 threads per block. *CSR1x1* comparatively results in poor thread granularity due to more number of blocks. *CSR1x1* results in an inefficient utilization as the multiprocessor spends a large fraction of time in block switching. Optimal block size (and grid size) is necessary to ensure maximal utilization of the resources.

With *NVIDIA's library*: The *CSR-vector* kernel uses one 32-thread warp per matrix row. This approach benefits from contiguous global memory access. However, as our testcases have small scale, the resulting matrices have fewer than 32 non-zeros per row, which causes underutilization of this kernel. Although this kernels outperforms our other SpMV considerations. A *CSR-scalar* which uses one thread per matrix row performs poorly compared to its vector counterpart. It uses 256 threads per block and outperforms *CSR1x1* which uses 16 threads per block. The *COO-vector* kernel is based on segmented reduction. The COO format has the worst computational intensity of all. The segmented reduction operation seems more expensive than alternative approaches that distribute work across threads of execution. *COO-scalar* approach which allocates only one thread has the worst computational complexity, equivalent of sequential CSR. Nevertheless, *COO-vector* kernel can be used to compensate the deficiencies of the other kernels. The *HYB* format couples the speed of ELL and the flexibility of COO. The ELL alone cannot be used in our matrices as there are very small number of nonzero entries per row. But when combined with *COO-vector* in *HYB kernel* we observe the performance improvement over *COO-vector* and *COO-scalar* kernels.

With *IBM's Library*: We experiment with two main compile-time optimizations proposed in [14], *CSR-Aligned Access*, aligned global memory accesses, and *Thread Mapping*, optimized thread mapping, for the CSR format. *Aligned Access* kernel is similar to *NVIDIA-vector* but uses 16-thread warp per matrix row and makes alignment adjustments. The matrices resulting from combination technique have very low number of nonzeros per row, usually less than 16, which contributes to the poor performance of these optimizations.

In terms of speed, the performance of the CPU remains constant (around 4 MFLOPs) while the performance of GPU based solvers boosts with the increasing problem sizes, see Figure 6 and Figure 7. CPUs generally mask the latency with on-chip cache, while GPUs mask the latency with a large number of threads and assigned work load. The increasing work load can promise to hide the latency to do texture accesses. With the problem sizes at our disposal, most of the cached

kernels fail to hide the latency causing unexpected slowdown. For small grids, both CPU and GPU implementations exhibit equivalent performance. Figure 8 displays a histogram which displays the distribution of number of grids as per the speedups achieved using GPU based solver. We observed that solvers achieve better speedups for the grids with large unknowns than for the smaller grids. For instance, with $d = 3$ and $n = 15$, one of the sparse grids has configuration of $N = 12285$ and $NNZ = 53229$, which achieves speedup of 98 using *CSR-Vector BiCGStab*. The cached versions also proved to be beneficial in this case.

VII. CONCLUSION AND FUTURE WORK

The sparse grid method is a practical tool to solve high dimensional PDEs numerically. In this paper we presented GPU based parallel implementations of Krylov subspace based iterative solvers for solving several small sized systems arising from this method. The efficiency and performance of various SpMV's and therefore of the linear solvers is demonstrated by numerical experimental results carried out on a NVIDIA Tesla device. Our results show that the choice of sparse format is not only important for scalability of iterative solvers to solve the sparse grids but the efficient implementation and parameter tuning of matrix-vector kernels is also essential for maximal performance on GPUs. The solution using grids with low approximation levels results in smaller problem sizes attributing to the overhead of the parallel platform and performs poorly compared to the sequential execution. We believe that with higher dimensions, the true potential of GPUs for high performance computing can be effectively tapped.

Further, in order to exploit the inherent parallelism in the combination technique, it is essential to distribute the solving of sparse grids over a cluster of GPUs. Such scheme poses several interesting future research directions. The variable times needed to solve the discrete systems would require either a reasonable distribution scheme or ways to decompose large systems over multiple GPUs. Future work will also include investigating the performance of these solvers on latest double precision NVIDIA devices.

REFERENCES

- [1] H. J. Bungartz, M. Griebel, D. R. Oschke, and C. Zenger. Pointwise convergence of the combination technique for the laplace equation. *East-West J. Numer. Math.*, 1994.
- [2] M. Griebel, T. Gerstner, and S. Wahl. *Method and device for evaluation of financial derivatives using sparse grids*. US Patent App. 09/994,114.
- [3] M. Griebel, M. Schneider, and C. Zenger. A combination technique for the solution of sparse grid problems. *Iterative Methods in Linear Algebra*, pages 263–281, 1992.
- [4] M. Griebel and V. Thurner. Efficient solution of fluid dynamics problems by the combination technique. *International Journal of Numerical Methods for Heat and Fluid Flow*, 5(3):251–269, 1995.
- [5] B. Lastdrager, B. Koren, and J. Verwer. Solution of time-dependent advection-diffusion problems with the sparse-grid combination technique and a rosenbrock solver. *Computational Methods in Applied Mathematics*, 1(1):86–98, 2001.
- [6] C. Pflaum. Convergence of the combination technique for Second-Order elliptic differential equations. *SIAM J. on Num. Ana.*, 34(6):2431–2455, Dec. 1997.
- [7] C. Pflaum and A. Zhou. Error analysis of the combination technique. *Numerische Mathematik*, 84(2):327–350, Dec. 1999.
- [8] J. Garcke, M. Griebel, and M. Thess. Data mining with sparse grids. *Computing*, 67(3):225–253, Oct. 2001.
- [9] O. Pironneau and Y. Achdou. *Computational Methods for Option Pricing*. Society for Industrial & Applied Mathematics, U.S., illustrated edition, 2008.
- [10] C. Reisinger. Analysis of linear difference schemes in the sparse grid combination technique. *SIAM Journal on Numerical Analysis*, 2006.
- [11] H. B. Zubair, C.C.W. Leentvaar and C.W. Oosterlee. Multigrid preconditioners for Bi-CGSTAB for the sparse grid solution of high-dimensional anisotropic diffusion equation. *J. Prime Res. Math.*, 3, 186-200 (2007)
- [12] C. Reisinger and G. Wittum. Efficient hierarchical approximation of high-dimensional option pricing problems. *SIAM J. on Scientific Computing*, 29(1):440, 2008.
- [13] L. Buatois, G. Caumon and B. Levy. Concurrent number cruncher - A GPU implementation of a general sparse linear solver. *Intl. Journal of Parallel, Emergent and Distributed Systems*, to appear.
- [14] M. Baskaran and R. Bordawekar. Optimizing Sparse Matrix-Vector Multiplication on GPUs. *IBM Technical Report RC24704.*, 2008.
- [15] J. Hull Options, Futures, and Other Derivative Securities *Prentice Hall, 4th Edition 2000.*, 2nd Edition, 1993.
- [16] T. Gerstner and M. Holtz. Valuation of performance-dependant options. *Applied Mathematics Finance*, 2008, 15, 1-20.
- [17] F. Heiss and V. Winschel. Likelihood approximation by numerical integration on sparse grids. *J. Econometrics*, 2008.
- [18] C. Zenger. Sparse Grids. In *Proceedings of the 6th GAMM seminar, Notes on numerical fluid mechanics*, volume 31, 1990.
- [19] A. Brandt. Multilevel adaptive solutions to boundary-value problems. *Math. Comp.*, 31:311329, 1977.
- [20] R. Barrett et al. Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods. *SIAM*, Philadelphia, 1994.
- [21] J. W. Demmel. Applied Numerical Linear Algebra. *SIAM*, Philadelphia, 1997.
- [22] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856-869, 1986.
- [23] J. J. Dongarra, I. S. Duff, D. C. Sorensen, and H. A. van der Vorst. Numerical Linear Algebra for High-Performance Computers. *SIAM*, Philadelphia, 1998.
- [24] H. A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of non-symmetric problems. *SIAM J. Sci. Stat. Comput.*, 13:631-645, 1992.
- [25] N. Bell and M. Garland Efficient Sparse Matrix-Vector Multiplication on CUDA. *NVIDIA Technical Report*, Dec. 2008.
- [26] S. Williams, L. Oliker, R. Vuduc and J. Shalf, K. Yelick, and J. Demmel. Optimization of sparse matrix-vector multiplication on emerging multicore platforms. *Proc. of the 2007 ACM/IEEE conf. Supercomputing*, 1-12, 2007.
- [27] J. Bolz, I. Farmer, E. Grinspun, and P. Schördor. Sparse matrix solvers on the GPU: conjugate gradients and multigrid. *ACM Trans. Graph.*, 22(3):917-924, July 2003.
- [28] J. Krüger and R. Westermann. Linear algebra operators for GPU implementation of numerical algorithms. *ACM Trans. on Graphics (TOG)*, 22(3):908-916, 2003.
- [29] V. Volkov and J. W. Demmel. Benchmarking GPUs to tune dense linear algebra. *Proc. of the 2008 ACM/IEEE conference on Supercomputing*, 2008, pp. 1-11.
- [30] N. Galoppo, N. K. Govindaraju, M. Henson, and D. Manocha. LU-GPU: Efficient algorithms for solving dense linear systems on graphics hardware. *Proc. of the 2005 ACM/IEEE conference on Supercomputing*, 2005.
- [31] G. Quintana-Orti, F. D. Igual, E. S. Quintana-Orti, and R. van de Geijn. Solving dense linear algebra problems on platforms with multiple hardware accelerators. *Proc. of the 14th ACM SIGPLAN Sym. PPPP*, 2009.
- [32] A. Cevahir, A. Nukada, and S. Matsuoka. Fast Conjugate Gradients with Multiple GPUs. *Computational Science, ICCS*, Vol.5544 Springer (2009), p.893-903.
- [33] D. Göddeke, R. Strzodka, J. Mohd-Yusof, P. McCormick, H. Wobker, C. Becker, and S. Türek. Using GPUs to Improve Multigrid Solver Performance on a Cluster. *Intl. Journal of CSE*, 2008.
- [34] http://www.nvidia.com/object/product_tesla_c1060_us.html