

HemeLB: A high performance parallel lattice-Boltzmann code for large scale fluid flow in complex geometries

M.D. Mazzeo, P.V. Coveney *

Centre for Computational Science, University College London, 20 Gordon St., London WC1H 0AJ, UK

Received 4 May 2007; received in revised form 4 December 2007; accepted 11 February 2008

Available online 4 March 2008

Abstract

We describe a parallel lattice-Boltzmann code for efficient simulation of fluid flow in complex geometries. The lattice-Boltzmann model and the structure of the code are discussed. The fluid solver is highly optimized and the resulting computational core is very fast. Furthermore, communication is minimized and the novel topology-aware domain decomposition technique is shown to be very effective for large systems, allowing us to tune code execution in geographically distributed cross-site simulations. The benchmarks presented indicate that very high performance can be achieved.

© 2008 Elsevier B.V. All rights reserved.

PACS: 47.00.00; 87.55.kd; 91.65.My

Keywords: Lattice-Boltzmann; Computer simulations; Parallel computing; Fluid dynamics; Complex systems

1. Introduction

The study of fluid flow in confined and complex geometries is of considerable interest in several domains including porous media [1] and the human vasculature [2,3]. Magnetic resonance imaging (MRI) [4] and advanced X-ray computed tomography (CT) [5] have been successfully applied to visualize and analyze flow in these situations. CT and magnetic resonance angiography (MRA) furnish non-invasive static and dynamic data acquisition [6]. In the clinical context, cerebral blood flow behavior plays a crucial role in the understanding, diagnosis and treatment of many conditions; indeed, some studies reveal relationships between specific flow patterns around vessel walls and cardiovascular diseases such as atherosclerosis [7]. However, experimental studies are often impractical owing to the difficulty of measuring transport properties in porous media and observing blood flow behavior (hemodynamics) in humans.

Modelling and simulation undoubtedly have a crucial role to play in these domains. The aforementioned imaging modalities provide high resolution data that simulation tools can exploit to accurately reproduce flow fields through the reconstructed geometries. Moreover, the need to capture fluid flow effects at multiple space scales with sufficient accuracy inevitably requires us to deal with large and complex systems, while many relevant processes are characterized by substantial time scales [8]. This problem is compounded when dealing with larger systems since computational fluid solvers take more time to converge. The greater number of simulation time steps required to achieve the desired behavior further increases the computational demand.

Modelling and simulation offer the prospect of providing clinicians with virtual patient specific analysis and thus diagnosis; moreover, they raise the possibility of performing non-invasive virtual experiments to plan and study the effects of certain courses of surgical treatment with no danger to the patient [2]. However, immediate clinical impact demands very rapid turn around of results, in close to real time if the methods are to be of interactive value.

* Corresponding author.

E-mail addresses: m.mazzeo@ucl.ac.uk (M.D. Mazzeo), p.v.coveney@ucl.ac.uk (P.V. Coveney).

Plainly, therefore, the availability of computational models of sufficient complexity and power is crucial. The computational fluid “solver” must itself be numerically highly efficient. While conventional continuum fluid solvers, based on finite difference, finite volume and finite element codes certainly exist [3,9], they are beset with problems in three spatial dimensions due to the computational costs of mesh generation, the need to solve the auxiliary Poisson equation for the pressure field, and various unavoidable approximations associated with the calculation of the shear stress from the flow velocity field. For large systems, such as those we are concerned with here, it is essential to develop and utilize scalable, high performance parallel codes, which further complicates the use of continuum models. In addition, the intricate geometry of porous media and blood vessels, as well as the treatment of fluid boundary conditions at such walls, are difficult for continuum fluid dynamics models to handle.

The lattice-Boltzmann method can be viewed as a special discretization of the continuous Boltzmann equation and is thus a kinetic approach that, unlike conventional numerical schemes, does not discretize the macroscopic continuum equations with the attendant need to solve the Poisson equation. Therefore, the lattice-Boltzmann (LB) method [10] offers an attractive alternative: our fluid solver is based on this modelling and simulation approach. Several papers have been published in hemodynamics¹ using the lattice-Boltzmann method [11–14].

The accuracy and performance of the lattice-Boltzmann method have been compared to those of finite-difference [15,16], finite-volume [17–19] and finite-element [19–22] methods. Various studies have confirmed that the lattice-Boltzmann method is competitive with the other approaches. Indeed, it is faster in situations where a specified accuracy is required, in particular in the context of the time-dependent simulation of large, complex systems by means of parallel implementations. However, a comparison between different fluid solvers is prone to ambiguity since their accuracy, intrinsic speed and convergence behavior all depend on the chosen parameters and specific details of the implementation.

The purpose of this paper is to introduce and report on a new lattice-Boltzmann model optimized for the study of fluid flow in complex sparse geometries. The paper is organized as follows. In Section 2, we review some related research. We describe the lattice-Boltzmann model in Section 3 and the boundary condition methods used in our code in Section 4. We then present the computational core and some single processor benchmarks on regular and complex systems in Section 5. The parallel implementation, associated benchmark results and a discussion are given in Section 6, followed by our conclusions in Section 7.

2. Overview of related work

Fluid flow simulation of very large and complex systems requires the use of a suitable physical model, substantial computational resources, as well as applications capable of exploiting them effectively. Much work has been done in order to speed-up single-processor and parallel lattice-Boltzmann simulations in regular systems. However, relatively few papers have been published which are aimed at the improvement of lattice-Boltzmann simulations of fluid flow confined in complex geometries. The parallel program presented in this work, called “HemeLB”, represents a contribution in this direction. We are currently using HemeLB to investigate cerebral blood flow behavior (“Heme” derives from hemodynamics) in patient-specific systems reconstructed via a graphical-editing tool we have developed.

In the context of LB fluid flow simulations in non-regular systems, particularly porous media [23–25], a number of studies have been carried out using the strategy presented by Donath et al. [26] to reduce the memory requirements. This approach avoids the storage of data associated with obstacles: 1D arrays, one for each data type, store the information about the fluid sites. The identification of the neighboring distribution functions, needed during the LB advection stage, relies on an extra 1D connectivity buffer. We note that an expensive global look-up table for the connectivity buffer is not required: if the 3D configuration file stores the data about all lattice sites it suffices to read this file only once in order to count the total number of fluid lattice sites, to allocate the 1D buffers and to build the connectivity by means of a look-up table composed of just two slabs of the entire system wherein the lattice site types and indices in the 1D buffers are stored. The slabs must iteratively scan the entire 3D domain and be swapped at every iteration. Thus, when the connectivity—supplied by an index or a pointer—needs to be established between a 1D buffer element and a neighboring (fluid) one, it is sufficient to find the data in the occupied slab.²

High performance lattice-Boltzmann simulations of large systems³ on cache-based microprocessors can be realized if there is sufficient spatial and/or temporal coherency since the poor bandwidth and latency of RAM memory cannot guarantee to keep the CPUs busy. This means that the data should be referenced frequently and/or with a sufficiently high level of sequentiality so that it is more likely to already reside in cache memory.⁴ The high cost of low latency memory limits the adoption of large ones. Usually,

¹ Hemodynamics is the research field which focuses on blood flow behavior.

² An approach that uses slabs to save memory is presented in detail in [27] in the context of the implementation of the lattice-Boltzmann method for regular systems.

³ A system is considered to be large if it does not fit in any of the caches.

⁴ While this is obvious for data accessed frequently, automatic prefetch of contiguous data guarantees their placement in cache before their processing. Some modern cache-based microprocessors prefetch non-contiguous data disposed in regular patterns and thus their workflow is optimized with respect to non-contiguous data arranged irregularly. The interested reader should consult Kowarschik and Weiß [28] for more detailed discussions on temporal and spatial coherency as well as organization of cache memories.

caches are organized in a hierarchical fashion in which the largest one, characterized by the lowest price per mm^2 , is the slowest and interfaces directly with RAM, while the smallest and fastest serves the CPU directly. This approach represents a good compromise in terms of cost and execution speed. Many cache-aware optimizations such as loop blocking, loop interchange, loop fusion and array merging [28] can substantially improve cache exploitation and thus reduce the execution time of many applications. Nevertheless, the presence of streaming and collision stages in lattice-Boltzmann simulations requires access to non-contiguous data, making it difficult to achieve high performance when simulating systems of even modest size. Pohl et al. [29] and Schulz et al. [30] presented two different “compressed grid” approaches which reduce the total memory consumption by a factor of almost two by exploiting the deterministic data dependencies which occur in the propagation step. The three-slabs approach of Martys and Hagedorn [31] and Argentini et al. [27] is also useful to reduce the total memory consumption. However, these memory-saving strategies do not solve the problem of modest performance on large systems since non-contiguous data needed during the propagation step often reside far from each other. 1D and 3D loop blocking techniques [29,32,33] show a similar performance improvement with respect to the approach of Martys and Hagedorn [31] and Argentini et al. [27] by augmenting space locality in small systems. The one technique which permits high performance to be attained on large systems seems to be the n -way blocking presented by Pohl et al. [29], Iglberger [32], and Velivelli and Bryden [34]. For 3D systems, 4-way blocking relies on maximum exploitation of the first-neighbor LB kernel by combining 3D space blocking with 1D time blocking. Many time steps are performed on 3D blocks of different shapes involving access to the whole domain at one time only. This strategy guarantees good spatial and temporal coherency and delivers high performance (see also Donath [35] and Kowarshik [36] for further benchmarks on cache-aware LB implementations). Relatively simple prefetch and preload techniques [37,38] can lead to a significant increase in memory bandwidth and improve the overall performance substantially. However, it is difficult if not impossible to apply the 4-way blocking approach to flow simulation in “sparse” (non-regular) systems because the data workflow is completely irregular.

Probably the best known data representation which is suited to LB fluid flow simulations in sparse and complex Cartesian geometries is that presented by Donath et al. [26] in the context of flow in porous media, where an approach based on regular 3D grids is abandoned in favor of 1D buffers which store the distribution functions of the fluid lattice sites and the connectivity information required during the propagation step. This strategy avoids the need to store data pertaining to non-fluid lattice sites. Sparse data representations have been successfully used in recent studies concerning flow in porous media [23,24].

However, while domain decomposition strategies based on cubes, slabs or parallelepipeds [39–44] yield good computational and communication load balancing when applied to regular systems, they are not well-suited to complex ones. Some domain decomposition strategies for lattice-Boltzmann codes which are well suited for non-regular lattices already exist, and we consider these briefly here. The state-of-the-art of domain decomposition, represented by the multilevel k -way partitioning scheme [45,46], can be used to partition non-regular systems. The need for users to explicitly implement these complex domain decomposition techniques can be avoided by using existing software such as the METIS library [47], which can be incorporated within parallel lattice-Boltzmann codes to obtain high quality domain decompositions and thus good execution performance on a large number of processors. However, even state-of-the-art parallel multi-level implementations require $O(N/10^6)$ seconds on $O(N)$ graph vertices⁵ even when performed on $O(100)$ processors [45]. This leads to unreasonable elapsed times when tackling systems with $O(10^8)$ or $O(10^9)$ fluid lattice sites. Moreover, these multilevel partitioning algorithms require a substantial amount of memory, which makes the use of a parallel approach compulsory as well as many processors even for systems which are not very large. Generally speaking, these multilevel methods use recursive bisection schemes for the initial partitioning. The interfaces between partitions are then calculated and load balancing is improved iteratively. As a consequence, these strategies have a large associated computational cost; unfortunately, they do not guarantee optimal communication or computational load balancing. For instance, the parallel 2D LB code presented by Dupuis and Chopard [48] employs the METIS library [47] to partition the system, but its performance on the small systems considered is poor. To compound this, large systems require long preprocessing times due to the high computational cost associated with the multilevel k -way partitioning method adopted in the METIS library [47].

The domain decomposition approach used by Pan et al. [49] is worthy of note. It combines a 1D data representation with connectivity [26] and is tested in Section 5.2; it is well-suited for sparse systems, possessing an efficient parallelization strategy based on the orthogonal recursive bisection (ORB) algorithm [50]. In the ORB approach, the computational grid is decomposed into two partitions such that the workload, due to fluid lattice sites only, is optimally balanced. The same process is then applied recursively $k - 1$ times to every partition, proceeding on the basis of orthogonal bisections. The total number of processors must be equal to 2^k . The amount of data required for an efficient implementation is proportional to the total number of fluid lattice sites N . Furthermore, the computational time scales as $O(N \log(N))$; while the domain decomposition is straightforward to implement and ensures good load balancing, it does not produce a satisfactory communication balance.

In the domain decomposition strategy employed by Wang et al. [51], as the system is read from file, N/P (P is the number of processors) fluid lattice sites are assigned to each processor consecutively. The difference of fluid lattice sites per processor is either 0 or 1, depending on whether N/P is an integer or not. Thus, the resulting computational load balance is perfect. The recovery of the data to be communicated is straightforward since the fluid lattice sites adjacent to neighboring processor sub-domains, called

⁵ In our case, a vertex in a graph corresponds to a fluid lattice site.




Fig. 1. Schematic diagram of a smooth boundary for which both horizontal inward-pointing directions indicated by the arrows originate from solid sites. Lattice sites within the computational fluid domain are denoted by empty circles. Many boundary condition methods are unable to determine the unknown distribution functions at the central fluid lattice site.

here “interface-dependent lattice sites”, are stored in an orderly fashion [51]. Hence the communication pattern is regular and no data flow discontinuities occur. The method is simple, assures perfect load balancing, does not require CPU time to optimize the decomposition and requires a memory consumption of $O(N/P)$ per processor, since global data are not necessary. Wang et al. [51] maintained their approach to be superior to that of the ORB technique. The domain decomposition approach itself is not new. In fact, it has been exploited in multilevel graph partitioning as an initial decomposition where it is referred to as the *as-is distribution*. However, such domain decomposition can be afflicted by poor communication balance even for simple systems, as will be shown in Section 6.1.

Finally, we note that space-filling curves [52] have been used to decompose sparse geometries in computational fluid dynamics (CFD) applications [53,54]. The quality and the speed of these partitioning strategies is sub-optimal, while the computational cost usually increases as $O(N \log(N))$, making them expensive for very large systems.

3. The HemeLB model

The lattice-Boltzmann model adopted in the HemeLB code is an incompressible version of the lattice Bhatnagar, Gross and Krook (BGK) D3Q15 one [55] (3D model with 15 velocity directions) proposed by Zou et al. [56], dubbed D3Q15i. The approximation of the continuum collision term is that suggested by Bhatnagar, Gross and Krook [57] for which the distribution function is assumed to evolve towards its local equilibrium value, at a rate controlled by a single relaxation parameter, τ :

$$\Omega \approx \frac{f^{(\text{eq})} - f}{\tau}. \quad (1)$$

The resulting lattice BGK (LBGK) equation is

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta x, t + \Delta t) - f_i(\mathbf{x}, t) = -\frac{f_i(\mathbf{x}, t) - f_i^{(\text{eq})}(\mathbf{x}, t)}{\tau}, \quad (2)$$

where the local equilibrium distribution functions are:

$$f_i^{(\text{eq})} = w_i \left(\rho + \frac{\mathbf{e}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{e}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2} \right), \quad (3)$$

where w_i is a weight coefficient, c_s is the speed of sound, \mathbf{e}_i is the velocity of the particle along the direction i and the hydrodynamic density ρ and macroscopic velocity \mathbf{u} are determined in terms of the distribution functions from

$$\rho = \sum_i f_i = \sum_i f_i^{(\text{eq})}, \quad \mathbf{u} = \sum_i \mathbf{e}_i f_i = \sum_i \mathbf{e}_i f_i^{(\text{eq})}. \quad (4)$$

The discrete velocities \mathbf{e}_i and the weight coefficients w_i must be chosen in order to ensure isotropic hydrodynamics (see Zou et al. [56] for details).

4. Boundary conditions

In this section, the boundary condition methods used in HemeLB are presented. Many boundary condition approaches have been investigated in the past, but few of them are well-suited to dealing with complex boundaries regardless of their orientation and shape. Probably the most studied and widely adopted wall boundary condition method is the bounce-back rule [58] for no-slip walls. Despite their number and variety, however, many boundary condition methods (e.g., [59–61]) are unable to handle the situation depicted in Fig. 1.

Instability issues and ambiguities, including those arising at boundaries where the normal and tangential components are not well defined, severely restrict the application of some boundary condition methods such as those used in [62] and [63]. The velocity boundary condition of Junk and Yang [64] handles the problem in Fig. 1, is second order accurate and thus in principle attractive.

However, it is very complex and computationally expensive, as it needs to store precise information related to the intersection between the boundary surface (which is not always available) and the lattice.

In the present work, we have implemented and studied various boundary condition methods. The non-equilibrium extrapolation method of Zhao-Li et al. [65] handles any type of boundary shape efficiently and is simple, making it worth investigating. By contrast the method of Zou and He [66] has very restricted applicability, but we have nevertheless tested it for non-inclined axis-aligned boundaries in situations where we can define the pressure (dubbed “pressure boundaries”).⁶ We have also developed and tested a new and simple extrapolation method for fluid flows confined by pressure and no-slip boundaries, which is outlined here. Let F be a no-slip boundary fluid lattice site. Before streaming, we assign the distribution functions at the new time step $\{f_i(F, t + \Delta t)\}$ to be the local equilibrium ones calculated with zero velocity and density employing the following extrapolation:

$$\rho(F, t + \Delta t) = a\rho(F, t) + b\rho(F, t - \Delta t), \quad (5)$$

with a, b real numbers such that $a + b = 1$. If F belongs to a pressure boundary, the extrapolated velocity is

$$\mathbf{u}(F, t + \Delta t) = a\mathbf{u}(F, t) + b\mathbf{u}(F, t - \Delta t). \quad (6)$$

When there are at least six unknown distribution functions at the inlet and outlet lattice sites, the method sets:

$$f_i(F, t + \Delta t) = f_i^{(\text{eq})}(\bar{\rho}, \mathbf{0}), \quad (7)$$

where $\bar{\rho}$ is the inlet or outlet pressure. The streaming must be performed from every non-solid lattice site. In our model, we chose $a = 1$ which reflects a zero-order extrapolation in time. This choice gives smaller errors than those provided using other boundary condition methods (as demonstrated at the end of this section) and the storage of densities and velocities at time step $t - \Delta t$ is not needed. Furthermore, the method is simple to implement and efficient considering that the local equilibrium distribution functions calculated at zero velocity require only a few multiplications and additions.

Finally, the resulting velocity and pressure imposed by velocity and pressure boundary conditions respectively are constrained to be equal to the prescribed ones and thus, for example, slip velocity cannot occur at a zero-velocity no-slip wall.

We first test our new boundary condition method by simulating fluid flow in a 3D square duct system subjected to a pressure gradient provided by an inlet pressure which is set to be greater than the outlet one. The analytical velocity profile is given by White [67]. The error analysis of the new extrapolation boundary condition method (denoted NewE) was carried out using the D3Q15i lattice Boltzmann model (see Section 3 for details) and compared with the bounce back rule (BB) and the non-equilibrium extrapolation method (NEE) presented by Zhao-Li et al. [65]. At the beginning of the simulation the distribution functions are set equal to the equilibrium one calculated with a unitary density and zero velocity, and steady-state is considered to be reached if

$$\frac{\sum_i \sum_j |u_x(i, j, t + \delta) - u_x(i, j, t)| + |u_y(i, j, t + \delta) - u_y(i, j, t)|}{\sum_i \sum_j |u_x(i, j, t)| + |u_y(i, j, t)|} \leq \delta Tol, \quad (8)$$

where $1 \leq i \leq nx$ and $1 \leq j \leq ny$ define the site location in every section perpendicular to the square duct axis, $lx \times ly \times lz = (nx - 1) \times (ny - 1) \times (nz - 1)$ is the system size in lattice units and $Tol = 10^{-8}$ with $\delta = 1/ly$. The maximum relative error is calculated to be

$$\text{err}_m = \frac{1}{u_0} \max \sqrt{(u'_x - u_x)^2 + (u'_y - u_y)^2 + (u'_z - u_z)^2}, \quad (9)$$

where (u'_x, u'_y, u'_z) is the analytical velocity and the maximum error is over the entire lattice [59,66]. The pressures at the inlet and outlet are regulated by means of the method of Zou and He [66] (dubbed “CoP”). In Table 1 we report the maximum relative error, the time steps required for convergence and the order of convergence, from a least-squares fit for selected Reynolds numbers and relaxation parameters (the last two parameters were chosen according to Zou and He [66]).

From Table 1, it emerges that our new boundary condition method is first-order accurate while its results are always more precise than the bounce back rule and the method of Zhao-Li et al. [65], except for the fastest flow in combination with the two smallest lattice sizes only, where the bounce back rule provides higher accuracy. From the method of Zhao-Li et al. the results are first-order accurate only (although they were asserted to be second order accurate in the original work [65]). Moreover, it is clear that our new method produces faster convergence rates than the others.

If we need to handle small systems with a specified accuracy, we can easily incorporate and use a more accurate boundary condition method such as the velocity boundary condition of Junk and Yang [64] if sub-grid-scale boundary information is provided.

⁶ Although this boundary condition method can only be applied to flat boundaries perpendicular to one of the three Cartesian axes, we decided to consider it because we found that its application for both the pressure boundaries and no-slip walls confining a rectangular duct yields results in accordance with 2D Poiseuille flow within machine accuracy, regardless of the Reynolds number, lattice resolution and viscosity. This excellent result is in contrast to that achieved by Zou and He [66] where such high accuracy was not reported.

Table 1

Maximum relative errors and orders of convergence, carried out with a least-squares fitting procedure, for 3D square duct flow modelled with various boundary condition (BC) methods (see Section 4 for their acronyms and references) and computational grid resolutions

LB parameters	BC method	$8 \times 4 \times 4$	$16 \times 8 \times 8$	$32 \times 16 \times 16$	$64 \times 32 \times 32$	Order
$Re = 10$ $\tau = 0.6$	CoP + NEE	3.963(−1) [417]	5.234(−1) [1959]	2.404(−1) [6842]	1.171(−1) [25834]	0.6399
	CoP + BB	3.071(−1) [697]	2.601(−1) [2345]	1.476(−1) [12140]	7.894(−2) [60707]	0.6697
	CoP + newE	7.239(−1) [224]	2.882(−1) [1266]	1.227(−1) [5676]	6.166(−2) [23154]	1.189
$Re = 5$ $\tau = 0.8$	CoP + NEE	1.134(0) [239]	7.893(−1) [778]	3.855(−1) [2616]	1.892(−1) [9446]	0.8784
	CoP + BB	4.128(−1) [325]	2.874(−1) [1545]	1.496(−1) [7553]	7.802(−2) [35603]	0.8153
	CoP + newE	2.688(−1) [130]	1.305(−1) [520]	6.062(−2) [2132]	3.425(−2) [8434]	1.002
$Re = 0.2$ $\tau = 1.1$	CoP + NEE	1.460(0) [148]	9.384(−1) [448]	5.284(−1) [1450]	2.805(−1) [5026]	0.7968
	CoP + BB	8.535(−1) [348]	3.840(−1) [1550]	1.724(−1) [6932]	8.320(−2) [30634]	1.123
	CoP + newE	3.557(−1) [105]	1.219(−1) [336]	5.151(−2) [1210]	2.723(−2) [4546]	1.236

The lattice size is $l_x \times l_y \times l_z$, the symbol $(-n)$ represents 10^{-n} and the numbers in square brackets are the iterations required to reach convergence.

Our method produces more stable simulations with respect to the other boundary condition methods. For $\tau = 0.6$, the simulation with the new boundary condition method is stable at $Re = 120$ while with the non-equilibrium extrapolation method and bounce-back instabilities occur at $Re = 100$ and $Re = 105$, respectively. Furthermore, the new scheme is very simple to implement and very efficient.

5. The computational core of HemeLB

In this section, we present a simple data structure which is adapted to sparse systems by avoiding retention of almost all information about the void component while achieving good data locality. The grid of the problem domain is split into a two-level representation: if any fluid lattice sites are present within a cell of the coarse grid, this is decomposed into a finer one. Such multi-level and hierarchical grids have previously been applied to lattice-Boltzmann simulations [19,62,68–70]. These approaches have been used in the context of grid refinement techniques to capture with sufficient accuracy small-scale hydrodynamic behavior within one or more sub-domains of the simulation domain. No data are exchanged between different grids belonging to the same level of resolution, but are passed instead to their parent coarser grids. In our approach, a two-level data structure is used to represent a domain with a single (unique) resolution. No data exchange takes place between the coarser (parent) grid and the finer one: the distribution functions are propagated within the latter only. This two-level structure saves memory with respect to the “full matrix” representation when one or more parent cells are completely occupied by non-fluid lattice sites since corresponding finer grids cannot then be allocated. Furthermore, interpreting each finer grid as a sub-domain for which good cache use takes place, spatial locality is automatically achieved as long as data located at the interfaces across different sub-domains, which we call “blocks”, are needed infrequently. For one single loop over the lattice sites of a block with side b in lattice units, the ratio between those interfacing with neighboring blocks and the interior ones scales as $1/b$. On the other hand, on increasing the block size cache utilization is likely to become worse, as occurs in medium and large regular systems.

Multi-level grids with more than two levels can be adopted. However, we found that a two-level grid with $b = 8$ represents a good tradeoff in terms of simplicity, cache use and adaptivity to sparse systems. With this choice, if the coordinate along the x axis of a lattice site is i , those of the corresponding block and its interior site are $bi = i \gg 3$ and $si = i - (bi \ll 3)$, respectively. Here, \ll and \gg shift the data to the left and right respectively; thus $n \ll 1 = n * 2$, $n \ll 1 = n * 4$ and so on. However, although this mapping between the coordinates of a lattice site in the single level representation and those in the two-level one is easy to implement, it is computationally expensive and should not be employed during the propagation step. We use instead a faster approach based on a look-up table of precomputed indices; this is described in the next section.

The 2D representation of our two-level hierarchical data organization is depicted in Fig. 2 where an idealized bifurcation is confined by a bounding box of 4×4 blocks of side $b = 4$ and discretized by the fluid voxels represented with solid edges. The data pertaining to the solid voxels in the white blocks are not stored and only the memory associated with the solid voxels represented with dashed edges is not used. We constructed a more compact data representation which only stores fluid voxel data in the preprocessing stage together with the allocation and initialization of the source and destination buffers of the distribution functions.

5.1. Serial benchmarks on regular systems

In the present subsection, timing results for standard lattice-Boltzmann implementations of regular systems are compared to those obtained with the two-level data structure presented in Section 5. The serial benchmarks have been carried out by adopting the D3Q19 model [55] for cubic systems of N^3 lattice sites. The bounce back rule is applied to the no-slip walls surrounding the systems. 3D and 4D data are stored in 1D arrays and a pre-computed lattice grid of integers `is_solid[N*N*N]` keeps track of




Fig. 2. 2D representation of our two-level hierarchical data organization for an idealized bifurcation with thick dashed edges. Its fluid lattice sites are depicted by means of voxels with solid edges. The bounding box comprises 4×4 blocks at the coarse level. Only the data associated with each non-solid block are allocated to fluid lattice sites as well as those solid ones represented with thin dashed edges.

the presence of the boundary conditions for each velocity direction l and lattice site n :

$$\text{is_solid}[n] \ \& \ (1 \ll l) = \begin{cases} 1 & \text{if its } l\text{th bit is 1 and } (i, j, k) + \text{dir}(l) \text{ is solid;} \\ 0 & \text{elsewhere.} \end{cases}$$

This procedure allows us to investigate simple problems while maintaining a level of generality related to our objective of addressing simulations of sparse systems. The implementations tested were not tuned to any specific hardware. Apart from the adoption of the 3D loop blocking technique in one implementation and an efficient way of referencing the lattice sites, no further optimizations have been implemented. Some data structures and optimizations have previously been investigated [29,32,33,35,37]. The code to handle the collision and propagation phases in one step only by means of the “collision-optimized” (or “array-of-structures”) data layout with the “push-scheme” [33] is illustrated in the following C lines (some declarations are omitted for brevity):

```
//source and destination buffers
//of distribution functions
double *old_f;
double *new_f;

old_f = (double *)malloc(sizeof(double) * N * N * N * 19);
...
// for each time step perform the following:
for (i = 0; i < N; i++) {
    for (j = 0; j < N; j++) {
        for (k = 0; k < N; k++) {
            Feq (&f_old[ijk = (i * N + j) * N + k], f_eq);

            for (l = 0; l < 19; l++) {
                ijk1 = ijk * 19 + l;
                f_old[ijk1] += omega * (f_old[ijk1] - f_eq[l]);
            }
            for (l = 0; l < 19; l++) {
                ijk1 = ijk * 19 + l;
                if (!(is_outside[ijk] & (1 << l)))
                {
                    f_new[(((i+e_x[l])*N+(j+e_y[l]))*N+(k+e_z[l]))*19+l] = f_old[ijk1];
                }
                else
                {
                    f_new[ijk * 19 + inv_dir[l]] = f_old[ijk1];
                }
            }
        }
    }
}
temp = old_f; // the buffers are swapped
old_f = new_f;
new_f = temp;
```

where `Feq ()` is the routine that calculates the local equilibrium distribution functions, ω is $-1/\tau$, `ex[l]`, `ey[l]`, and `ex[l]` are the three components of the lattice direction `l`, and `inv[l]` indicates the opposite direction to `l`. We refer to this implementation as the “standard” one. We note that many operations can be avoided in the standard implementation. For example, `((i+e_x[l])*N+...)*19+1` depends only on the reference lattice site⁷, `N` and `l`, and has the form `ijk + f_inc[l]` where `f_inc[]` is a precalculated integer array of 19 elements. Exploiting this array and avoiding unnecessary additions and multiplications as far as possible, an optimized standard code is produced. This is called “Opt Standard” and has the following structure:

```
double *f_old_a, *f_old_b;
double *f_new_p;

int *is_out_p;
...
// for each time step perform the following:
f_old_a = f_old;
f_new_p = f_new;
is_out_p = is_outside;

for (i = 0; i < N*N*N; i++) {
    f_old_b = f_old_a;
    lbmFeq (f_old_a, f_eq);

    for (l = 0; l < 19; l++)
    {
        *f_old_a += omega * (*f_old_a - f_eq[ l ]); ++f_old_a;
    }

    for (l = 0; l < 19; l++)
    {
        if (!(*is_out_p & (1 << l)))
        {
            *(f_new_p + f_inc[ l ]) = *f_old_b;
        }
        else
        {
            *(f_new_p + inv_dir[ l ]) = *f_old_b;
        }
        ++f_old_b;
    }
    ++is_out_p;
    f_new_p += 19;
}
temp = old_f;    // the buffers are swapped
old_f = new_f;
new_f = temp;
```

Our new method for handling sparse geometries efficiently is based on a two-level data structure that can be represented by a 2D array instead of a 1D array, with one dimension for the block identifier and one to identify the lattice site and the velocity direction. For the purpose of comparison with respect to the optimized standard implementation, the arrays used are again chosen to be 1D. In this way, the collision and propagation steps are identical and the principal factor that makes the difference in performance is the data reordering due to the two-level representation. However, a strategy is needed to identify efficiently the neighboring lattice sites involved in propagation. We observe that each lattice site has a neighboring one whose parent block coordinates differ from the reference one by -1 , 0 or 1 . Furthermore, the identifier of the neighboring lattice site in its block always ranges between 0 and $8^3 - 1$.

At this point, a neighbor map can be constructed based on the foregoing considerations. Since the data are stored in 1D arrays, a 1D lookup table `f_map[]` of $8^3 \times 19$ elements can incorporate all the information regarding the increment in coordinates associated with the neighboring block and the neighboring lattice identifier. The lookup table is constructed in the following manner:

⁷ The access to an nD array in C requires n additions. For example, `ex[l]` is equivalent to `*(ex + l)`.


```

// declarations
m = -1;
for (i = 0; i < 8; i++) {
  for (j = 0; j < 8; j++) {
    for (k = 0; k < 8; k++) {
      for (l = 0; l < 19; l++) {
        ni = i + e_x[l];
        nj = j + e_y[l];
        nk = k + e_z[l];
        si = ni - ((bi = ni >> 3) << 3);
        sj = nj - ((bj = nj >> 3) << 3);
        sk = nk - ((bk = nk >> 3) << 3);
        block_inc = (bi * (N >> 3) + bj) * (N >> 3) + bk;
        site_id = (((si << 3) + sj) << 3) + sk;
        f_map[ ++m ] = (block_inc * (8*8*8) + site_id) * 19 + l;
      }
    }
  }
}

```

where (i, j, k) and (ni, nj, nk) are the coordinates of the reference lattice site and those of the neighboring one in the single level representation, respectively, (si, sj, sk) are the coordinates of the latter in the two-level representation and (bi, bj, bk) the increments in coordinates of the neighboring block with respect to the reference one. The resulting implementation, called “newA ($8 \times 8 \times 8$)”, is based on `f_map[]`. If the two-level representation is based on blocks of 2^3 lattice sites, the lattice site identifiers are reordered as in the Morton scheme [71]. The resulting implementation is denoted “newA ($2 \times 2 \times 2$)”. The array `is_solid[]` is still exploited to handle no-slip boundaries, but it has a two level representation in this case.

We note that the array `is_solid[]` is not needed as long as the neighboring lookup table is extended to identify the neighboring lattice sites for all velocity directions, at the expense of significant memory consumption. The neighboring lookup table is represented in two-level fashion and the resulting full connectivity implementations are dubbed “newB” and “newB+”, respectively, according to whether the effect of the inclusion of the “IF” branch related to the bounce back rule is considered or not. In “newB” the indices are set equal to the key value “−1” whenever the corresponding neighboring locations reside in the wall. An implementation of the 3D blocking technique with loop blocking of 8, denoted “blocking ($8 \times 8 \times 8$)” has been applied to the optimized standard implementation in order to increase space coherence, cache reuse and thus performance for the simulation of large systems. All the implementations tested have undergone the same level of optimization. Usually, a halo region of lattice sites surrounding the system handles the bounce back rule more efficiently than the “IF” branch present in the propagation step. However, fusion of streaming and collision steps, fast referencing and an optimized implementation of the equilibrium distribution functions guarantee rapid executions here.

Apart from the possibility of tuning the codes presented in this section to the specific hardware used, the version of HemeLB which implements the 3D blocking technique represents a state-of-the-art LB implementation which is well-suited to regular systems when a predefined order involving the propagation step, as in the compressed-grid and 4-way blocking algorithms [29,32], is not exploited. In any case, these methods cannot be applied to sparse systems where data blocking yields better performance than those achieved using mesh reordering schemes [26].

Our benchmarks for the LB implementations were conducted using an Opteron 2 GHz with a L2 cache of 1024 KB, Linux kernel 2.6.19 and an IBM Power4 1.7 GHz with 1.9 MB L2 cache and IBM AIX OS. The compiler used on the Opteron is the Intel version 8.1 with flags `-O3 -unroll -unroll-loops -IPF-fma -axW -xW -inline` while on the Power4 (RISC architecture) the executable was produced with `xlc -qstrict -qstrict_induction -O5 -qipa=level=2 -qhot=vector -qcache=auto -lm -bmaxdata:0x70000000`. Double precision is used to calculate local equilibrium distribution functions and their moments. The performance as a function of the cube root of the total number of lattice sites is measured in millions of site updates per second (MSUPS), a standard performance measure for lattice Boltzmann codes [26,27,29,32–35,37,38,51], and is shown in Fig. 3 for the codes presented in this section.

The superiority of implementations which adopt the two-level data structure with respect to the standard one (and its optimized version) is evident on the Opteron machine. On the Power4, the large cache and extensive hardware support for prefetch, branch prediction, and both out-of-order and speculative execution of instructions optimize cache use and execution on large systems. Thus, the performance degradation for larger systems is low and the simple optimized standard implementation performs very well. The implementations “newA ($8 \times 8 \times 8$)” and “newB ($8 \times 8 \times 8$)” exhibit similar performance on both processors and are often the fastest executions. The two-level implementation with full connectivity “newB+ ($8 \times 8 \times 8$)” is slightly faster than “newB ($8 \times 8 \times 8$)” due to the elimination of the expensive “IF” branch in the propagation step while the standard implementation is substantially inferior to its optimized counterpart. We conclude that redundant operations cannot always be avoided by modern compilers and should be eliminated by the programmer.

On the Opteron, the 3D loop blocking implementation exhibits significantly higher performance than its non-blocking counterpart, indicative of good cache reuse. Two-level data representations deliver very high performance. It is interesting, however, that the

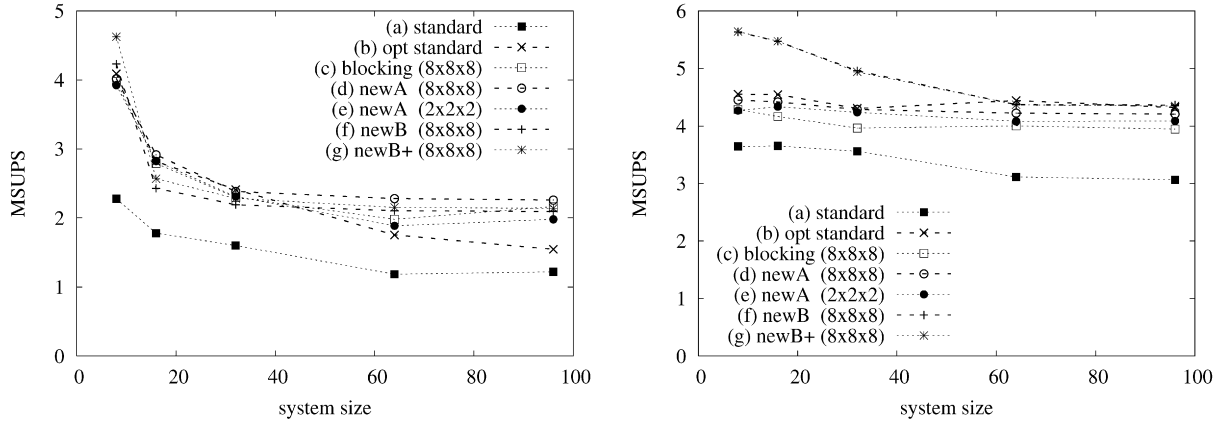


Fig. 3. Serial performance measured in millions of lattice site updates per second (MSUPS) as a function of the cube root of the total number of lattice sites for (a) the standard (“standard”); (b) its optimized counterpart (“opt standard”) and (c) its 3D loop blocking version with factor 8 blocking (“blocking ($8 \times 8 \times 8$)”); that using the two-level representation with (d) a look-up table of indices used in the propagation step with blocks of 8^3 (“newA ($8 \times 8 \times 8$)”) and (e) 2^3 (“newA ($2 \times 2 \times 2$)”) lattice sites, respectively; (f) a two-level buffer of indices to handle the propagation step (“newB ($8 \times 8 \times 8$)”); (g) as in (f) supplying the bounce back rule without the “IF” branch (“newB+ ($8 \times 8 \times 8$)”). The timing results were conducted on an Optron 2 GHz (left) and a Power4 1.7 GHz (right) (see Section 5.1 for further details).




Fig. 4. Volume rendering of the velocity flow field in a bifurcation of 6253409 fluid lattice sites characterized by $\tau = 0.6$ and a peak velocity of 0.066 lattice units. The wall and pressure boundary conditions applied to the inlet and outlet lattice sites and no-slip walls, respectively, are discussed in Section 4. Yellow and blue colors represent maximum and minimum (zero) speeds, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

two-level data structure with 2^3 blocks (Morton) does not produce the same benefit as an implementation with larger blocks. Overall, we have demonstrated that data reordering techniques can produce higher performance than typical cache-aware optimizations such as loop blocking.

5.2. Benchmarks on sparse systems

Here, we compare the performance of the two-level data layout with that achieved by the standard approach normally used for LB simulations of fluid flow confined in sparse geometries as presented by Donath et al. [26]. In the standard implementation, a single level representation is constructed by scanning respectively over the x -, y -, and z -directions, while in our new strategy a two-level representation is constructed following the allocated fluid blocks directly.

We carried out the benchmarks by applying the incompressible D3Q15i model (see Section 3) to a set of bifurcations of different sizes. The simulation cells are composed of 16^3 , 32^3 , 64^3 , 128^3 and 256^3 lattice sites with 225, 1601, 12449, 98211 and 780359 fluid lattice sites, respectively. The volume rendering of the velocity field of the largest system (using $\tau = 0.6$ and characterized by a peak velocity of 0.2) is shown in Fig. 4. The collision and propagation steps are implemented as in the piece of C code labeled “newB” in Section 5.1. The boundary conditions implemented are those presented in Section 4. During streaming the “IF” branch is avoided at boundary lattice sites by allocating $N * 15 + 1$ values (N is the number of fluid lattice sites and 15 the number of




Fig. 5. Serial performance measured in millions of site updates per second (MSUPS) as a function of the cube root of the fluid lattice sites for the standard (“standard”) and the new approach described in Section 5.2. The timing results were conducted on an Opteron 2 GHz (left) and a Power4 1.7 GHz (right) (see Section 5.2 for further details).

velocity directions) for the source and destination buffers: the last (ghost) value of the destination buffer $f_new[]$ is referenced when a particle streams towards the wall.

The processors, operating systems and compilers are those used in Section 5.1 for the benchmarks on cubic systems. The performance in MSUPS is shown in Fig. 5 for the standard and new approaches described in this section.

We first note the excellent performance attained with both kinds of processors. Our new approach exhibits higher performance on the two largest systems using the Opteron and very similar behavior in the other cases. As anticipated, the two-level-grid-based approach yields better performance on large systems than the standard technique because data are likely to reside more closely in memory address space and thus in cache. The ratio between wall fluid sites and others is greater for smaller systems. Since the wall fluid sites are handled with less operations than the others (see Section 4), the aforementioned ratio as well as cache use affect the performance behavior as a function of system size.

The performance of our lattice Boltzmann algorithms for more complex systems has been investigated by means of a parallel implementation, which is presented in Section 6.

6. Parallel implementation

In this section we first present the parallelization strategy adopted in HemeLB (Section 6.1) which is shown to be well-suited to simulate fluid flow in large, complex systems. Then, our new partitioning strategy is described and compared with other ones in terms of advantages and disadvantages. The new domain decomposition approach is found to be very fast and of high quality (Sections 6.1 and 6.3). The computational performance achieved with the adoption of the domain partitioning technique presented here is reported in Section 6.3.

6.1. Domain decomposition

Parallel lattice Boltzmann codes well-suited for complex systems often employ graph partitioning techniques such as the multilevel k -way partitioning scheme [30,48] or the ORB approach [23,24,49,72] to decompose the computational domain into a set of high quality subdomains so as to achieve good workload and communication balancing during the parallel execution. Here, a domain decomposition approach based on a graph growing partitioning (GGP) algorithm, a specific branch of the general category of graph partitioning techniques, is presented. GGP partitioning techniques have been presented and used in the past [73–75], but have never been adopted in LB simulations. The GGP scheme does not scan the system in a predefined order and, as a consequence, it is afflicted by one of the main limitations of many other domain decomposition strategies: the need to store global data for efficient execution. Moreover, while the computational cost is optimal with respect to the number of elements to distribute it remains $O(N/10^6)$ seconds on $O(N)$ elements. Nevertheless, the nature of the GGP is likely to yield good communication and computational load balance. In this section, the GGP algorithm is reviewed in the context of LB simulations of sparse systems and then a novel modification is presented which overcomes limitations related to the global data requirements and inherent speed. Finally, the important impact of the proposed GGP scheme on the output is outlined.

First, we describe the basic GGP method. Basically, the GGP technique grows P partitions from P corresponding seed lattice sites employing a breadth-first expansion algorithm [23,24,49,72]. At the beginning of the GGP scheme, all the fluid lattice sites are marked as “non-visited”. A fluid lattice site is selected and marked as “visited” and assigned to the processor with identifier (or rank) 0. An iterative search starts from this lattice site and grows a region around it until the desired number of (visited) fluid lattice

sites N/P is reached, as explained below. The expansion proceeds by means of two “coordinate” buffers that are used to locate the fluid lattice sites at the current and next iteration, respectively. For example, at the first iteration, the first coordinate buffer will point to the first selected visited fluid lattice site only, while the second one will contain the coordinates of the non-visited neighboring fluid lattice sites. At the end of every iteration the two aforementioned buffers are swapped. The propagation from every fluid lattice site in the first coordinate buffer to the neighboring non-visited ones is constrained to follow the velocity directions in the LB model used. When the number of fluid lattice sites reaches N/P , a set of $P - 1$ similar iterative searches, one for each processor, starts from other non-visited fluid lattice sites until the total number of fluid lattice sites visited is N . Particular care must be taken if an iterative search cannot continue because the visited fluid lattice sites in the first coordinate buffer are completely surrounded by solid lattice sites but the number of visited fluid lattice sites has not yet reached N/P . In this case, another iterative search from a non-visited fluid lattice site is started.

Next, we outline some implementation and performance issues. We note that, in general, each fluid lattice site is checked more than once: the iterative search must start from a non-visited lattice site which is not known *a priori*; the expansion does not proceed in a predefined order and, inevitably, the iterative search checks many fluid lattice sites to see if they are available (not visited). A dynamic look-up table of the non-visited fluid lattice sites available at every iteration, and the velocity directions from which the visited fluid lattice sites originate, speeds up the iterative search by minimizing repeated access at the expense of an increase in memory consumption. However, the algorithm updates $O(10^6)$ fluid lattice sites per second on modern processors and, consequently, the domain decomposition requires only a few minutes for a system of hundreds of millions of fluid lattice sites. Furthermore, the memory consumption is $O(N)$ bytes.

Our improvement on the naive GGP algorithm relies on the two-level data representation previously presented (Section 5) and is described here. Basically, the GGP method is only applied to the coarser grid of the two-level data hierarchy. The input file is read a first time to store the number of fluid lattice sites for every block of the coarsest grid, a step which can be avoided if already accomplished in a preprocessing stage. The resulting lookup table, dubbed `fluid_sites_per_block[]`, is global but has a very low memory overhead.⁸ The GGP algorithm is then applied to the coarse grid only, as mentioned before. Here, all the fluid lattice sites of a block are assigned to a processor.

This coarse-level-based GGP approach plays a crucial role in managing the overall computational and memory costs as well as in the output of the results. First of all, the computational cost is reduced by three orders of magnitude with respect to the original GGP approach if each block is composed of $O(1000)$ lattice sites.⁹ This allows us to decompose a system of $O(10^9)$ fluid lattice sites in a second. Moreover, each processor will need to store detailed information on the lattice sites contained only in the blocks marked with its identifier (rank) and their neighboring ones. This is accomplished by reading the input file a second time and using `proc_id[]`. In this way, the memory consumption is kept of $O(N/P)$ plus a small contribution due to the global lookup tables `fluid_sites_per_block[]` and `proc_id[]`.¹⁰ 2D representations of the result of the proposed domain decomposition method applied to a square and a bifurcation with 128² and 61451 fluid lattice sites, respectively, are provided on the right-hand side of Fig. 6, while Wang et al.’s method [51] produced the images on the left hand side. There are always 16 partitions and each block of the finer resolution of the two-level representation contains 8² lattice sites. The grey-scale helps one to recognize the identifier of each partition. The two methods always assign 1024 lattice sites to each partition in the square system and this results in perfect load balancing. However, the shape of the partitions influences the communication balance. With our new partitioning method the minimum and maximum number of interface-dependent fluid lattice sites, i.e. the fluid sites adjacent to neighboring processor sub-domains, are halved and thus inter-processor communications are reduced substantially. In the bifurcation, the smallest and largest partitions produced with the new technique contain 3439 and 3792 fluid lattice sites respectively for which the load balance would be sub-optimal. While Wang et al.’s method [51] ensures perfect load balancing, its interface-dependent fluid lattice sites in the smallest and largest partitions are respectively 190/212 and 2370/876 of those attained with our new method.¹¹

The parallel LB scheme presented by Schulz et al. [30] can overlap the inter-processor non-blocking communications with the computation of the interface-independent fluid lattice sites. This may result in considerable benefit since the communications may be partially or completely hidden. We found that the standard vendor MPI implementation (vMPI) used to handle inter-processor communications does not overlap non-blocking communications with computation regardless of the amount of (perfect) computational balance imposed. As a consequence, achieving good computational and communication balance turns out to be crucial. Moreover, a faster computational core makes the effect of large and unbalanced communications even more critical [39, 51]. From Fig. 6, we can say that Wang et al.’s method [51] is a generalization of the slab-based partitioning technique for sparse systems and inherits the latter’s communication imbalance even when P is very low with respect to N . In contrast, the success of GGP techniques, including that presented here, resides in the nearly-spherical propagation of the visited fluid sites at each

⁸ For a simulation cell of 1024³ fluid sites the memory requirement is around 4 MB only, assuming that every block contains 8³ lattice sites and 2 bytes are needed for every element of `fluid_sites_per_block[]`.

⁹ This was the choice we adopted since every block contains 8³ = 512 lattice sites.

¹⁰ Eventually, the memory due to `fluid_sites_per_block[]` can be freed before any local $O(N/P)$ memory allocation.

¹¹ The differences in the values are also due to the presence of system boundaries and hence, the inner partitions are more likely to share a larger number of distribution functions with the neighboring ones.




Fig. 6. 2D partitions of a square and a bifurcation with 128^2 and 61451 fluid lattice sites respectively, as obtained with the domain decomposition method of Wang et al. [51] (left-hand side images) and that presented in Section 6.1 with a block size of 8^2 for the finest resolution of the two-level representation (right-hand side images); see Section 5 for details. In each image the number of partitions is 16 and their ranks are colored according to a grey-scale map. The identifiers of the interface-dependent distribution functions between every pair of adjacent partitions are represented by tiny grey segments. The grey intensity is proportional to the identifier. Adjacent distribution functions have the same identifier and cannot be communicated during each LB time step.

iteration.¹² However, perfect load balancing of the sort that characterizes the original GGP approach is no longer guaranteed, since the search proceeds for arbitrary blocks of lattice sites, for which a sub-optimal load balancing is always ensured.

6.2. Optimizations

We now discuss a variety of optimizations incorporated in HemeLB. The resulting code is very fast yet hides the considerable complexity inherent in simulations of sparse systems, minimizing the amount of intra-machine communication with respect to a specific domain decomposition while optimizing inter-machine communication in cross-site runs in grid deployments during each LB time step, keeping low those communications needed for the output of flow fields.

6.2.1. Minimization of communication

Sparse and complex systems present various difficulties in terms of parallel implementations. Generally, following domain decomposition, the identifiers of the distribution functions to be communicated and the location of the corresponding fluid lattice sites must be determined since they are not known *a priori*, as would be the case for regular domain decompositions. One parallel approach which does not need this step is that of Wang et al. [51]. Communicating the identifiers doubles¹³ the size of inter-processing communications at each time step. In HemeLB this is avoided; instead, only the distribution function values must be

¹² In general, the precise shape of the advancing front involved in the GGP search depends on the graph to be partitioned, the presence of obstacles and of visited graph vertices. At each iteration of the GGP search, the visited fluid lattice sites involved at the previous step are arranged in a cube for the D3Q15, D3Q18, D3Q19 and D3Q27 LB models if the propagation has not encountered any obstacle.

¹³ Usually, each distribution function is stored in 8 bytes and its velocity direction and the location of the interface-dependent fluid lattice site cannot be stored in only 4 bytes for large systems.

communicated. After domain decomposition is achieved, each processor calculates its interface-dependent identifiers. Without loss of generality, let us consider only two neighboring processors A and B. A and B have precomputed a buffer of identifiers, labeled $f_id[]$.¹⁴ In general, $f_id[m]$ calculated by processor A is not equal to that of B for any m . This non-matching of identifiers is overcome by performing one single point-to-point communication from A to B if A is smaller than B. In this way, the buffer $f_id[]$ of B is overwritten with that of A. This procedure needs to be applied to every neighboring processor in the preprocessing phase only. During each LB time step, only the values of the interface-dependent distribution functions need be communicated.

6.2.2. Transparency of implementation complexity

We have designed HemeLB's computational core based on the optimizations presented in Section 5.1 and further ones discussed in this section. The optimizations discussed in this section, Section 5 and the previous one permit us to maintain a low overall memory consumption and reduce the execution time by a factor of three. The implementation complexity pertaining to sparse geometries is a source of several possible optimizations. As we have seen in Section 5.1, removing redundant operations and complex access to buffer elements can reduce the execution time considerably. The optimization presented in Section 6.2.1 represents another improvement in this direction because the use of extra buffers of indices and many associated operations is avoided.¹⁵ We stress that, in HemeLB, two-level representations (see Section 5) are represented by 1D arrays. Thus, the access to an element of an array requires minimal computational cost. The complexity associated with sparse systems is hidden in extra connectivity buffers. One of them is responsible for the on-processor propagation of particles between fluid lattice sites, as we have seen in Section 5.

Now we describe how we have optimized the set-up of the buffers to be copied due to communications between neighboring processor sub-domains. The streaming of particles between two different processor sub-domains takes place in three steps: (a) copying of the distribution functions from the source buffer $f_old[]$ to that to send to the neighboring processor $f_to_send[]$; (b) copying of the latter buffer into the receiving one $f_to_recv[]$ by means of a point-to-point communication; (c) copying of the distribution functions from the receiving buffer $f_to_recv[]$ to the destination one $f_new[]$. Steps (a) and (c) represent the difficulty that must be faced if one wants to avoid the use of extra buffers corresponding to multiple indices, because the interface-dependent fluid lattice sites are not arranged regularly and are not known *a priori*. The problem is solved by using two extra buffers of indices with 4 bytes per element: $f_old_id[]$ and $f_new_id[]$ are set during the preprocessing step (following domain decomposition) through the buffers $f_id[]$ and a two-level local lookup table $f_map[]$ that maps from the lattice locations to the source and destination buffer indices. In this way, steps (a) and (c) are simply accomplished with the following lines of code

```
f_to_send[ : ] = f_old[ f_old_id[:] ];
...
f_new[ f_new_id[:] ] = f_to_recv[ : ];
```

where the symbol `:` means that all the elements of the arrays must be referenced. In this way, we have also optimized the copying of the distribution functions to be communicated between neighboring processor sub-domains.

Next, we describe a trick to avoid insertion of a computationally intensive conditional in the inner loop related to the propagation step. At first, the streaming towards void or fluid lattice sites covered by neighboring processor sub-domains, which must be handled by communications, appears to be avoided only by including a conditional indicating whether the neighboring lattice site is contained on an adjacent processor sub-domain or not. However, the “IF” branch may be eliminated by streaming the distribution functions in the source buffer, which would propagate towards neighboring processor sub-domains or solid sites, to a non-used ghost element of the destination one. If the processor sub-domain has N/P fluid lattice sites, the ghost buffer element is covered by the $[(N/P) * 15 + 1]$ th one. The source buffer must also be sized with $(N/P) * 15 + 1$ elements, since at the end of the LB time step they must be swapped.

With the adoption of the features presented in this section, the resulting code becomes very simple as the programming complexity caused by sparse representations is then completely hidden, avoiding complicated buffer element referencing and costly conditionals.

6.2.3. Optimization of inter-machine communication

The exploitation of multiple machines permits us to tackle “grand-challenge” problems including very-large-scale fluid flow simulations. Indeed, many problems in biological and physical sciences require computational resources beyond any single supercomputer's current capacity [77–80]. The advent of cross-site simulations is already taking place today by means of various geographically distributed MPI interfaces, such as MPICH-G2 [81] and its new version MPIg [82]. A first investigation of the feasibility and scalability of cross-site HemeLB simulations using geographically distributed domain decomposition (GD^3) [80] is presented in Section 6.3. The communication between geometrically remote machines is characterized by a latency one or two orders of magnitude greater than that affecting intra-machine data exchanges while the inter-machine bandwidth is considerably

¹⁴ In the HemeLB program, each element contains 8 bytes, 6 for the global lattice location and 2 for the velocity direction involved in the communication.

¹⁵ Access to a n -dimensional array involves n additions in C. The situation is worse in programming languages like Fortran where many additions and multiplications are required.

lower than the intra-machine one. Thus, the optimization of inter-machine communications is essential for such an approach to be viable.

The GGP technique produces partitions of high quality when the ratio between the total number of fluid lattice sites and the number of processors used is large. When this ratio is small, the system can be firstly subdivided into m parts with m small ($O(10)$), then each part into further m parts, and so on. This approach is equivalent to the ORB method (see Section 2) if $m = 2$ and the partitions are obtained proceeding from orthogonal bisection. Thus, if the number of machines used is m , the strategy adopted is to subdivide the system into m partitions. The size of each one is chosen to be proportional to the number of processors of the corresponding machine. Even though the topology is described in the Globus RSL (resource specification language) job description file [83], the application is unable to access this file to discover the topology. Another method is required to inform the application about the topology; two possible methods are (a) to create an extra input file to specify the topology or (b) to use MPIg routines to discover it. The latter is the preferred option as it does not require maintenance of two distinct files. Thus, we have implemented a topology discovery mechanism based on the MPI function `MPI_Attr_get` in HemeLB in order to detect the number of machines used and the number of processors for each machine. Furthermore, the inter-machine non-blocking communications are split out from the others, then hidden within the computation of the other communications and the computation of interface-independent fluid lattice sites whenever possible, following the scheme of Schulz et al. [30]. Briefly, in this parallelization strategy, after the collision of the interface-dependent fluid lattice sites the propagation of their particles to the neighboring processors is initiated by means of non-blocking communications. Completion of these communications is enforced only after the computation of interface-independent fluid lattice sites. The topology discovery implementation, the two-level GGP decomposition and the adoption of the Schulz et al. parallelization scheme [30] permit us to (a) optimize inter-machine communications so as to hide high latency and low memory bandwidth and (b) tune cross-site simulations to specific heterogeneous resource distributions. Unfortunately, among MPIg, MPICH-G2 and vMPI, only MPIg hides inter-machine communications via non-blocking communications that occur within the supercomputing resources and the MPI implementations used in our benchmarks. The current parallel implementation of HemeLB does not optimize the domain decomposition to take into account the different performance of the plurality of multiprocessor machines that may be present in a grid. However, this heterogeneity may be taken into account simply by adopting the algorithmic scheme employed to consider the different numbers of processors used in each machine as discussed above.

6.2.4. Attainment of pattern regularity

Except for the partitioning scheme presented in Wang et al. [51], in which the interface-dependent fluid lattice sites are stored contiguously, more sophisticated domain decomposition techniques are prone to yield pattern irregularity since the data to be communicated are, in general, not located in close proximity within memory address space. This induces poor cache use when handling those data and engenders lower overall performance. The use of the extra connectivity buffer for the propagation step makes no difference to the precise order of the indices of the source and destination buffers. Hence, pattern regularity of the interface-dependent fluid lattice site data is achieved by locating them closely in the last part of the source and destination buffers. A second benefit stems from the use of direct memory addressing of the inner fluid lattice sites, as explained below. The adoption of the parallel scheme [30] relies on the subdivision of interface-dependent fluid lattice sites from the other ones. In the original work, this classification was handled by additional expensive buffers and costly indirect addressing operations. Here, instead, inner fluid lattice sites are stored contiguously and thus the corresponding elements in the source buffer are accessed sequentially, resulting in significant memory and computational gains.

6.2.5. Output optimization

In this section, we outline the code optimization we have performed regarding simulation output. The output data are chosen to be the effective pressure, velocity and von Mises stress flow fields in single precision. Therefore, the memory consumption is $15 \times 8 / (5 \times 4) = 6$ times lower than that obtained by storing all the distribution functions for every lattice site.¹⁶

Checkpointing¹⁷ is essential (a) to avoid the need to restart the simulation from the beginning if it crashes and (b) to monitor results interactively and modify the configuration parameters during the course of the simulation (“steering” [84]). Full checkpoint capabilities have been integrated in HemeLB and the portable binary format used, XDR (external data representation), allows different platforms to work with the same input and output data.

With the adoption of the GGP domain decomposition technique presented in Section 6.1, the time to write the flow fields to output can be reduced substantially with respect to many other partition approaches if data must be presented in a specific order. In this case, assuming that only the processor 0 writes to the output file but cannot store the partial flow fields of all the other processors, $O(N)$ single point-to-point communications must be performed since, for example, the first fluid site may belong to processor 3 say, the second one to the processor 1 say, and so on. This problem is minimized in axis-aligned domain decomposition approaches where the partitions are regular; it is wholly absent in the method of Wang et al. [51] since the data are subdivided in an orderly fashion among the processors.

¹⁶ It is assumed that each fluid lattice site has 15 distribution functions stored in double precision.

¹⁷ During checkpointing the entire system state is dumped at every user-defined simulation period.




Fig. 7. Snapshot of the volume rendering of the velocity flow field pertaining to a patient-specific intracranial vasculature system comprising 7.7 million fluid lattice sites. Yellow and blue colors represent maximum and minimum (zero) speeds, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The GGP domain decomposition technique presented in Section 6.1 proceeds for blocks of m^3 lattice sites. As a consequence, the single point-to-point communications drop from $O(N)$ to $O(N/m^3)$, that is by around three orders of magnitude if $m = 8$. The resulting communication cost becomes reasonable since the effects due to high latency and limited bandwidth are kept low.

6.3. Performance on single and distributed multiprocessor architectures

In this section, we report some benchmarks for the parallel HemeLB code presented in Section 6. Single-site parallel performance has been tested on HPCx, the IBM SP p690+ Power4 1.7 GHz housed at Daresbury Laboratory (UK) [85], on Bigben, the Cray XT3 MPP TeraGrid machine located at the Pittsburgh Supercomputing Center (USA) [86] and on HECToR, a Cray XT4 located at the University of Edinburgh's Advanced Computing Facility [87]. The IBM machine is composed of a number of nodes communicating via the IBM High Performance Federation Switch. Every node is comprised of 8 chips, each of 2 processors that share L2 and L3 caches of 1.9 and 36 MB, respectively. The dual-core AMD Opteron 2.6 GHz CPUs of Bigben platform run the Catamount Operating System (OS), while the front end processors run SuSe Linux. The interconnect and storage network are Cray SeaStar and Infiniband respectively. The GNU compiler 4.0.2 was used with flag `-O3`. The dual-core AMD Opteron 2.8 GHz CPUs of HECToR run the UNICOS/Linux OS. The communications network utilizes Cray SeaStar2 communication chips. The Portland Group compiler was used with flags `-Minform=severe -O3 -fastsse -DXT3 -Mipa=fast,inline -Munroll=c:16 -Munroll=n:8`.

Benchmarks were conducted on two square ducts of $256 \times 128 \times 128$ and $1024 \times 512 \times 512$ fluid lattice sites respectively, a bifurcation of 6.25 million fluid lattice sites (see Fig. 4) and two patient-specific systems of 4.69 and 7.70 million fluid lattice sites respectively obtained from medical MRA datasets reconstructed by means of a graphical-editing tool we have developed (Fig. 7 shows volume rendering of the velocity flow field of the biggest patient-specific system). The bounce back rule and the boundary condition method presented in Section 4 were used to tackle the regular systems and the complex ones, respectively. The minimum number of processors used for the $1024 \times 512 \times 512$ system was 128. We compared the performance delivered using the new GGP method (see Section 6.1) to that achieved with two parallel versions in which (a) an ideal domain decomposition approach partitions the smallest system and (b) the Wang et al. method [51] is adopted when tackling the bifurcation. HPCx was exploited for this purpose. We applied the most favorable configuration to the latter partitioning method: the bifurcation is oriented along the horizontal x -axis and the partitioning scheme applied by scanning the z - and y -directions before x ; this produces partitions with fewer interface-dependent fluid lattice sites than that attained by cutting the bifurcation into horizontal subdomains.

In Fig. 8 we report the single-site HemeLB performance in MSUPS as a function of the processor count on the various machines and for the regular and complex systems described above. All benchmarks were obtained without checking for instabilities or applying convergence criteria. However, their impact is very modest since they involve only one “All-reduce” communication of a few bytes which is performed at every n time steps where n may be large ($O(100)$).

We note that in general our LB fluid solver achieves excellent performance. The ratio of fluid lattice sites between adjacent subdomains to the interior ones is smaller for larger systems; this effect amortizes the communications cost and, as a consequence, the parallel scalability behavior tends to be higher. HemeLB yields outstanding performance when simulating systems with complex geometries for which the scalability and intrinsic performance in terms of MSUPS often exceed those attained in the other cases. Our GGP scheme operates upon the coarser level of the two-level data hierarchy characterizing HemeLB (see Sections 5 and 6.1).

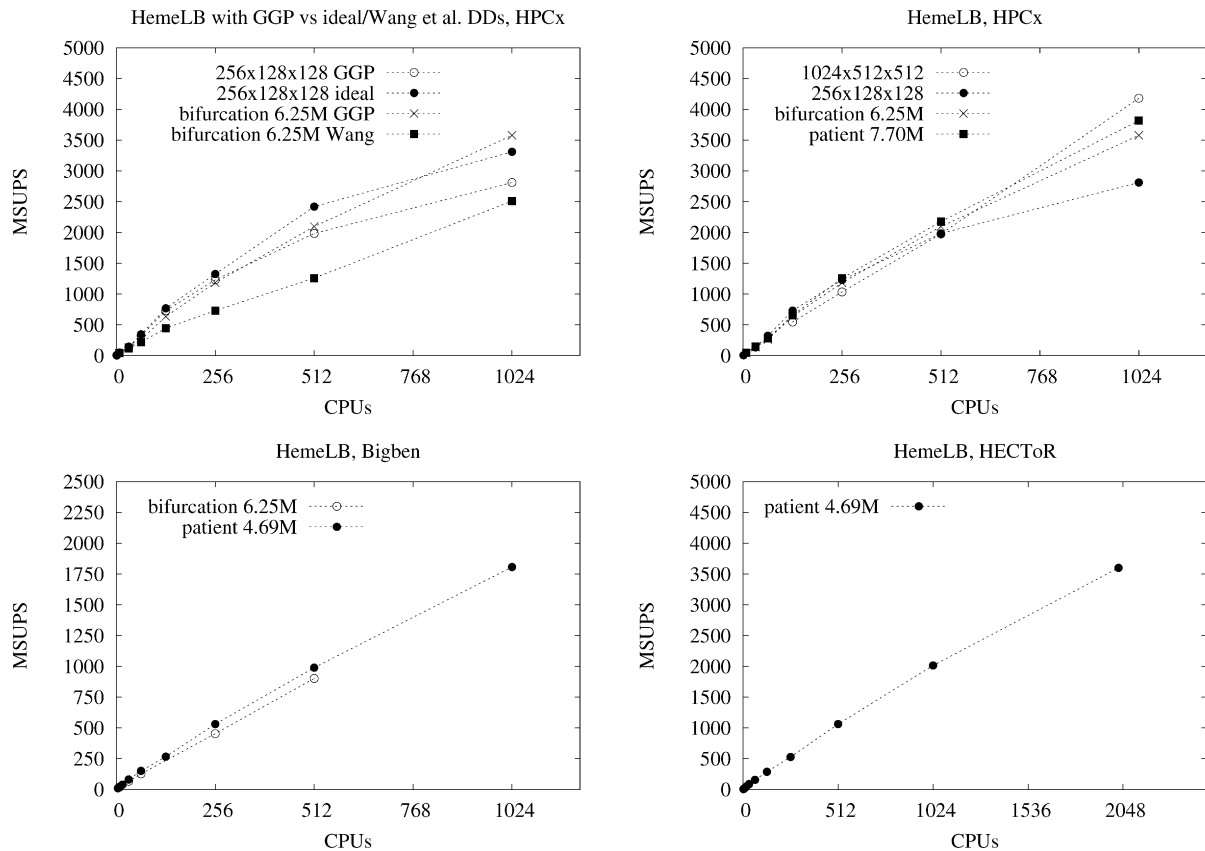


Fig. 8. HemeLB single-site parallel performance measured in millions of lattice site updates per second (MSUPS) as a function of the number of the CPUs used for the regular systems of $256 \times 128 \times 128$ and $1024 \times 512 \times 512$ fluid lattice sites, a large bifurcation of 6.25 million fluid lattice sites and two patient specific systems of 4.69 and 7.7 million fluid lattice sites, respectively, on various platforms (see Section 6.3 for more details). Apart from the GGP method presented in Section 6.1, denoted “GGP”, on HPCx we also used an ideal regular partitioning scheme and the domain decomposition approach presented by Wang et al. [51] in the parallel code to tackle the smallest regular system and the bifurcation, respectively.

We stress that, except for the largest regular system, the number of fluid lattice sites in each subdomain is comparable to that in a macrocell of the coarse grid when the processor count is 512 or more. Thus each subdomain comprises a few of those macrocells. Despite the presence of this very fine granularity, our GGP technique ensures good load balancing and high parallel scalability. All these results are demonstrated in our benchmarks. The smallest regular system is executed at a rate of 175 time steps per second on 128 processors. Except for the highest processor count, we note that HemeLB with our new GGP-based partitioning technique produces similar performance to that achieved by adopting a regular domain decomposition. With 1024 CPUs the communication cost dominates over the computation. For example, analyzing again the benchmark associated with the smallest regular system, the communication and the computational cost per time step associated with interface-independent lattice sites calculated by the processor with rank 0 are 0.9515 and 0.3738 ms, respectively. In this connection, we remark that the use of an MPI implementation which overlaps communication with computation is essential in order to attain very high performance, especially when N/CPUs is low.

The GGP technique plays an important role in the outputting of results, as anticipated in Section 6.1 and 6.2.5. For example, the output results are written in 12.9 and 11.59 seconds using 1 and 1024 processors, respectively. Thus, the point-to-point communications between the processor that writes to disk and the other ones have negligible impact. On the largest regular system, $\text{MSUPS}(\text{CPUs} = 1024) = 2.12\text{MSUPS}(512)$ indicative of superlinear speedup, due to the high quality of the partitioning carried out with our GGP technique and better cache use when the processor count is higher, i.e., when each subdomain is smaller.

The time required to accomplish the domain decomposition calculated by the processor with rank 0 and that required to manage the rest of the preprocessing phase without considering input reading—buffer management (BM)—are 0.482 and 3.336 seconds on 128 processors and 0.4483 and 0.391 seconds 1024 processors, respectively. As already mentioned, our GGP scheme can domain-decompose very large systems and is exceptionally fast because it operates on a coarse grid. Furthermore, BM is reasonably fast and scales very well. Our new domain decomposition approach divides the bifurcation into sub-domains of higher quality than those obtained by the Wang et al. method [51]. For example, the latter technique produces several sub-domains with no interface-independent fluid lattice sites when the processor number is 512 or 1024. This is due to an insufficient number of available vertical




Fig. 9. HemeLB cross-site performance in millions of lattice site updates per second (MSUPS) as a function of the number of LONI clusters, each contributing 32 processors (see Section 6.3). Ideal performance is depicted with a dashed line. MPIg was employed to handle inter-machine communications.

layers to decompose the system into partitions with thickness greater than one in lattice units.¹⁸ The performance superiority of our domain decomposition approach over the other implementations is evident for a number of CPUs greater than or equal to 32. When the processor count is 1024, communication of distribution functions between different sub-domains and computation associated with interface-independent lattice sites, calculated by the processor with rank 0, are 0.7559 and 0.7374 ms per time step, respectively. In spite of this, we believe that the performance could be enhanced by a factor of two if we had available an MPI implementation which could overlap communication with computation.¹⁹ On Bigben the performance is around four times lower than that attained on 1024 processors of HPCx but is characterized by an excellent scalability. For example, $\text{MSUPS}(\text{CPUs} = 512) = 1.988 \text{ MSUPS}(256)$ for the simulation of the bifurcation.

Cross-site benchmarks were performed on the IBM P5-575 AIX clusters housed at Louisiana Tech, Tulane University and ULL, called Bluedawg, Ducky and Zeke respectively comprising part of LONI (Louisiana Optical Network Infrastructure). Each node of each cluster comprises 8 IBM Power5 1.9 GHz processors. The network interface is an IBM High Performance (Federation) Switch [88]. The compiler flag used is -O2 only because we experienced some problems using very high compiler optimizations.

We tested HemeLB on the largest patient-specific system (7.7 million fluid lattice sites). In Fig. 9, the performance in MSUPS as a function of the number of LONI clusters used is provided; 32 processors were employed for each cluster and the grid-enabled MPI interface adopted for inter-machine communications is MPIg. We note high scalability and parallel efficiency even with the employment of all three clusters which is an impressive result considering that the simulation is running across geographically distributed platforms.

It is difficult to compare HemeLB's performance with that of other parallel lattice-Boltzmann codes developed for complex systems. Each code has been tested using different lattice-Boltzmann models, architectures and flow field systems. However, to our knowledge, the parallel implementation that yields the best performance results in terms of scalability and overall execution speed is that presented by Wang et al. [51]. We have directly demonstrated that the domain decomposition strategy adopted in [51] is inferior to that incorporated in HemeLB. Moreover, HemeLB exploits several novel optimizations that significantly speed up computation and the output of flow fields. Furthermore, we may mention that among the best parallel codes for regular systems (referenced in Section 2), only those presented in [44,76] are comparable to HemeLB in terms of speed-up and fluid site updates per second for a particular processor count using non-vector machines. This is an excellent result for HemeLB since our code has not been tuned to any specific architecture, while in [44,76] specific hardware optimizations and regular domain decomposition approaches were applied. Finally, in none of our work to date have we sought the most favorable compiler for HemeLB and its available optimizations. We discovered *a posteriori* that these aspects are very important and may substantially improve single core execution.

7. Conclusions

We have described a very efficient parallel lattice-Boltzmann code which is well-suited for application to sparse fluid systems. The fast computational core relies on a two-level data representation. This methodology guarantees the minimization of memory consumption when tackling complex sparse systems and yields a faster execution than conventional implementations due to better

¹⁸ Each side of the minimum bounding box that contains the whole system is smaller than 512 in lattice units.

¹⁹ In this case only the communication and interface-dependent computational costs would affect the simulation while the computation due to interface-independent lattice sites would be accomplished before completion of communication.

cache use. Several optimizations are described that reduce redundant operations, increase pattern regularity, simplify the computational core, reduce memory consumption and optimize intra-machine communications. The novel topology-aware two-level domain decomposition is very fast and guarantees high quality domain decomposition partitions. We have demonstrated its capability to effectively decompose a very large system in less than half a second, and its superior communication balance compared to other methods previously used in flow simulations in porous media, while ensuring good workload distribution. Consequently, the parallel scalability and performance of the HemeLB code are very good. The ability to overlap non-blocking communications with computation whenever possible through the use of an optimized MPI implementation is crucial if maximum computational performance is to be realized. Furthermore, the domain decomposition adopted has a direct and positive effect on the output of the flow fields, drastically reducing point-to-point communications. Moreover, a simple modification of the domain decomposition technique permits us to optimize inter-machine communications in grid-based cross-site runs. As a consequence, HemeLB can effectively tackle large and complex systems via distributed heterogeneous computational resources. We note that HemeLB is well-suited to efficiently simulate regular systems. In conclusion, HemeLB delivers very high performance on multiprocessor machines for both regular and complex (sparse) fluid flow systems.

One of our major aims is to use HemeLB for the study of a variety of human cerebral blood flow scenarios, ranging from normal to neuropathological conditions, including aneurysms and arterio-venous malformations, as well as the entire intra-cranial vasculature. The excellent multiprocessor performance attained using heterogeneous resources is particularly attractive for addressing patient-specific cerebral blood flows in clinically relevant wallclock times. The use of sufficient aggregated computational power should enable us to tackle large systems at interactive rates. We plan to develop HemeLB as part of a complete problem solving environment wherein the medical data from various imaging modalities are manipulated by graphical editing tools, embedded within HemeLB, and rendered by state-of-the-art visualization software. The full checkpoint capabilities within HemeLB allow simulation parameters to be changed during the course of a run and output data to be monitored through steering tools. These will be integrated in the near future in order to effectively investigate time-dependent simulation flows and to create a robust and direct connection between the application that manipulates the medical data, HemeLB, and the tool which visualizes the output flow fields.

Acknowledgements

We thank Dr. Shawn Brown at Pittsburgh Supercomputing Center and Kevin Roy of Cray Research for conducting benchmarks on Cray XT3 and XT4 platforms, receptively. We would like to thank Luis Fazendeiro of the Centre for Computational Science (CCS), University College London (UCL), Prof. Frank Smith and Dr. Nick Ovenden of the Mathematical Modeling Group (MMG) at UCL for helpful discussion on boundary conditions; also Daniel Scott, Stefan Zasada, Dr. Radhika Saksena, and Dr. Simon Clifford of the CCS for their generous general technical support. We are indebted to Dr. Brian Toonen and Prof. Nicholas Karonis of the Mathematics and Computer Science Division at Argonne National Laboratory as well as TeraGrid support for technical assistance and help in setting up cross-site runs. We are grateful to Dr Steven Manos (CCS) for producing the data exhibited in Fig. 9. The benchmarks on TeraGrid machines were obtained under NFS grants DMR070013 and DMR070014. MDM is grateful to EPSRC for funding his studentship in association with Reality Grid (GR/67699/02 and EP/C536451/01). This work was also partially funded by EPSRC under the GENIUS project (EP/F00561X/1).

References

- [1] S. Chen, D. Doolen, Lattice Boltzmann method for fluid flows, *Annu. Rev. Fluid Mech.* 30 (1998) 329.
- [2] C.A. Taylor, M.T. Draney, J.P. Ku, D. Parker, B.N. Steele, K.C. Wang, C.K. Zarins, Predictive medicine: Computational techniques in therapeutic decision-making, *Comput. Aided Surgery* 4 (1999) 231.
- [3] R. Botnar, G. Rappitsch, M.B. Sheidegger, D. Liepsch, K. Perktold, P. Boesiger, Hemodynamics in the carotid artery bifurcation: a comparison between numerical simulations and in vitro MRI measurements, *J. Biomech.* 33 (2000) 137.
- [4] M.D. Mantle, A.J. Sederman, L.F. Gladden, Single- and two-phase flow in fixed-bed reactors: MRI flow visualisation and lattice-Boltzmann simulations, *Chem. Eng. Sci.* 56 (2001) 523.
- [5] Y. Nakashima, Y. Watanabe, Estimate of transport properties of porous media by microfocus X-ray computed tomography and random walk simulation, *Water Resour. Res.* 38 (2002) 1272.
- [6] M. Goyen, M.E. Ladd, J.F. Debatin, J. Barkhausen, K.H. Truemmler, S. Bosk, S.G. Ruehm, Dynamic 3D MR angiography of the pulmonary arteries in under four seconds, *J. Magn. Reson. Imag.* 13 (2001) 372.
- [7] M.J. Thubrikar, F. Robicsek, Pressure-induced arterial-wall stress and atherosclerosis, *Annals of Thoracic Surgery* 59 (1995) 1594.
- [8] G. Giupponi, J. Harting, P.V. Coveney, Emergence of rheological properties in lattice Boltzmann simulations of gyroid mesophases, *Europhys. Lett.* 73 (2006) 533.
- [9] K. Imaeda, F. Goodman, Analysis of non-linear pulsatile flow in arteries, *J. Biomech.* 13 (1980) 1007.
- [10] S. Succi, *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond*, Oxford University Press, 2001.
- [11] H. Fang, Z. Wang, Z. Lin, M. Liu, Lattice Boltzmann method for simulating the viscous flow in large distensible blood vessels, *Phys. Rev. E* 65 (2002) 051925.
- [12] M. Hirabayashi, M. Ohta, D.A. Rufenacht, B. Chopard, Characterization of flow reduction properties in an aneurysm due to a stent, *Phys. Rev. E* 68 (2003) 021918.
- [13] A.M. Artoli, A.G. Hoekstra, P.M.A. Sloot, Mesoscopic simulations of systolic flow in the human abdominal aorta, *J. Biomech.* 39 (2006) 873.
- [14] R. Ouared, B. Chopard, Lattice Boltzmann simulations of blood flow: Non-Newtonian rheology and clotting processes, *J. Stat. Phys.* 121 (2005) 209.

- [15] D.R. Noble, J.G. Georgiadis, R.O. Buchius, Comparison of accuracy and performance for lattice Boltzmann and finite difference simulations of steady viscous flow, *Int. J. Num. Meth. Fluids* 23 (2003) 1.
- [16] K. Sankaranarayanan, I.G. Kevrekidis, S. Sundaresan, J. Lu, G. Tryggvason, A comparative study of lattice Boltzmann and front-tracking finite-difference methods for bubble simulations, *Int. J. Multiphase Flow* 29 (2003) 109.
- [17] M. Breuer, J. Bernsdorf, T. Zeiser, F. Durst, Accurate computations of the laminar flow past a square cylinder based on two different methods: Lattice-Boltzmann and finite-volume, *Int. J. Heat Fluid Flow* 21 (2000) 186.
- [18] J. Bernsdorf, F. Durst, M. Schäfer, Comparison of cellular automata and finite volume techniques for simulation of incompressible flows in complex geometries, *Int. J. Numer. Meth. Fluids* 29 (1999) 251.
- [19] S. Geller, M. Krafczyk, J. Tölke, S. Turek, J. Hron, Benchmark computations based on Lattice Boltzmann, Finite Element and Finite Volume Methods for laminar flows, *Ergebnisberichte des Instituts für Angewandte Mathematik*, Nummer, 2004, p. 274.
- [20] D. Kandhai, D.J.-E. Vidal, A.G. Hoekstra, H. Hoefsloot, R. Iedema, P.M.A. Soot, Lattice-Boltzmann and finite element simulations of fluid flow in a SMRX static mixer reactor, *Int. J. Num. Meth. Fluids* 31 (1999) 1019.
- [21] D.O. Martínez, W.H. Matthaeus, S. Chen, D.C. Montgomery, Comparison of spectral method and lattice Boltzmann simulations of two-dimensional hydrodynamics, *Phys. Fluids* 6 (1994) 1285.
- [22] S. Chen, Z. Wang, X. Shan, G.D. Doolen, Lattice Boltzmann computational fluid dynamics in three dimensions, *Phys. Fluids* 68 (1992) 379.
- [23] X.X. Zhang, A.G. Bengough, J.W. Crawford, I.M. Young, A lattice BGK model for advection and anisotropic dispersion equation, *Adv. Water Resour.* 25 (2002) 1.
- [24] X.X. Zhang, A.G. Bengough, L.K. Deeks, J.W. Crawford, I.M. Young, A novel three-dimensional lattice Boltzmann model for solute transport in variably saturated porous media, *Water Resour. Res.* 38 (2002) 1167.
- [25] X.X. Zhang, L.K. Deeks, A.G. Bengough, J.W. Crawford, L.M. Young, Determination of soil hydraulic conductivity with the lattice Boltzmann method and soil thin-section technique, *J. Hydrology* 306 (2005) 59.
- [26] S. Donath, T. Zeiser, G. Hager, J. Habich, G. Wellein, Optimizing performance of the lattice Boltzmann method for complex structures on cache-based architectures, in: F. Huelsemann, M. Kowarschik, U. Ruede (Eds.), *Frontiers in Simulation: Simulation Techniques—18th Symposium in Erlangen*, September 2005 (ASIM), SCS Publishing House, Erlangen, 2005, p. 728.
- [27] R. Argenti, A.F. Bakker, C.P. Lowe, Efficiently using memory in lattice Boltzmann simulations, *Future Gener. Comput. Syst.* 20 (2004) 973.
- [28] M. Kowarschik, C. Weiß, An Overview of Cache Optimization Techniques and Cache-Aware Numerical Algorithms, in: *Lecture Notes in Computer Science*, vol. 2625, Springer, 2002, p. 213.
- [29] T. Pohl, M. Kowarschik, J. Wike, K. Iglberger, U. Rude, Optimization and profiling of the cache performance of parallel lattice Boltzmann codes in 2D and 3D, *J. Parallel Process. Lett.* 13 (2003) 549.
- [30] M. Schulz, M. Krafczyk, J. Tolke, E. Rank, Parallelization strategies and efficiency of CFD computations in complex geometries using Lattice Boltzmann methods on high-performance computers, in: M. Breuer, F. Durst, C. Zenger (Eds.), *High-Performance Scientific and Engineering Computing, Proceedings of the 3rd International FORTWIHR Conference on HPSEC*, Erlangen, 2001, Springer-Verlag, Berlin, 2002, p. 115.
- [31] N.S. Martys, J.G. Hagedorn, Multiscale modeling of fluid transport in heterogeneous materials using discrete Boltzmann methods, *Mater. Struct.* 35 (2002) 650.
- [32] K. Iglberger, Cache optimizations for the lattice Boltzmann method in 3D, Bachelor Thesis, Lehrstuhl für Informatik 10 (Systemsimulation), Institut für Informatik, University of Erlangen-Nuremberg, Germany, 2003.
- [33] G. Wellein, T. Zeiser, G. Hager, S. Donath, On the single processor performance of simple lattice Boltzmann kernels, *Comput. Fluids* 35 (2006) 910.
- [34] A.C. Velivelli, K.M. Bryden, A cache-efficient implementation of the lattice Boltzmann method for the two-dimensional diffusion equation, *Concurrency Comput.: Pract. Exper.* 16 (2004) 1415.
- [35] S. Donath, On optimized implementations of the lattice Boltzmann method on contemporary high performance architectures, Lehrstuhl für Informatik 10, Institut für Informatik, University of Erlangen-Nuremberg, Bachelor thesis, 2004.
- [36] M. Kowarschik, Data locality optimizations for iterative numerical algorithms and cellular automata on hierarchical memory architectures, PhD thesis, Lehrstuhl für Informatik 10, Institut für Informatik, Universität Erlangen-Nürnberg, ISBN 3-936150-39-7, 2004.
- [37] S. Hausmann, Optimization and performance analysis of the lattice Boltzmann method on x86-64 based architectures, Bachelor Thesis, Erlangen, 2005.
- [38] A.C. Velivelli, K.M. Bryden, Optimizing performance on modern HPC systems: Learning from simple kernel benchmarks, in: *Proceedings of the 2nd Russian-German Advanced Research Workshop on Computational Science and High Performance Computing*, HLRS, Stuttgart, 2005.
- [39] J. Ni, J. Zhang, C.-L. Lin, S. Wang, Parallelization of a lattice Boltzmann method for lid-driven cavity flow, in: *UI research booth, Supercomputing*, 2003.
- [40] J.-C. Desplat, I. Pagonabarraga, P. Bladon, LUDWIG: A parallel lattice-Boltzmann code for complex fluids, *Comput. Phys. Comm.* 134 (2001) 273.
- [41] T. Pohl, F. Deserno, N. Thurey, U. Rude, P. Lammers, G. Wellein, T. Zeiser, Performance evaluation of parallel large-scale lattice Boltzmann applications on three supercomputing architectures, in: *Proceedings of the ACM/IEEE SC2004 Conference, Supercomputing*, 2004, p. 21.
- [42] P. Giovanni, F. Massaioli, S. Succi, High-resolution lattice Boltzmann computing on the IBM SP1 scalable parallel computer, *Comput. Phys.* 8 (1994) 705.
- [43] G. Amati, S. Succi, R. Piva, Massively parallel lattice-Boltzmann simulation of turbulent channel flow, *Int. J. Mod. Phys. C* 8 (1997) 869.
- [44] G. Wellein, P. Lammers, G. Hager, S. Donath, T. Zeiser, Towards optimal performance for lattice Boltzmann applications on terascale computers, in: A. Deane, et al. (Eds.), *Parallel Computational Fluid Dynamics—Theory and Applications. Proceedings of the Parallel CFD 2005 Conference*, ISBN 0-444-52206-9, College Park, MD, USA, May 24–27, 2005, Elsevier, 2006, p. 3.
- [45] G. Karypis, V. Kumar, Parallel multilevel k -way partitioning scheme for irregular graphs, *SIAM Rev.* 41 (1999) 279.
- [46] G. Karypis, V. Kumar, Multilevel k -way partitioning scheme for irregular graphs, *J. Parallel Distrib. Comput.* 48 (1998) 96.
- [47] G. Karypis, K. Schloegel, V. Kumar, Parallel graph partitioning and sparse matrix ordering library, <http://www-users.cs.umn.edu/karypis/metis/parmetis/files/manual.pdf>, 2003.
- [48] A. Dupuis, B. Chopard, Lattice gas: An efficient and reusable parallel library based on a graph partitioning technique, in: *HPCN99*, in: *Lecture Notes in Comput. Sci.*, vol. 1529, Springer, 1999, p. 319.
- [49] C. Pan, J.F. Prins, C.T. Miller, A high-performance lattice Boltzmann implementation to model flow in porous media, *Comput. Phys. Comm.* 158 (2004) 89.
- [50] V. Kumar, A. Grama, G. Karypis, *Parallel Computing*, The Benjamin/Cummings Publishing Company Inc., California, 2004.
- [51] J. Wang, X. Zhang, A.G. Bengough, J.W. Crawford, Domain-decomposition method for parallel lattice Boltzmann simulation of incompressible flow in porous media, *Phys. Rev. E* 72 (2005) 016706.
- [52] H. Sagan, *Space-Filling Curves*, Springer-Verlag, 1994.
- [53] M.J. Aftosmis, M.J. Berger, G. Adomavicius, A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries, in: *38th AIAA Aerospace Sciences Meeting and Exhibit*, Reno NV, Jan. 2000.

- [54] M.J. Aftosmis, M.J. Berger, S.M. Murman, Applications of space-filling curves to Cartesian methods for CFD, in: 42nd AIAA Aerospace Sciences Meeting and Exhibit, Jan. 2004.
- [55] Y. Qian, D. d'Humières, P. Lallemand, Recovery of Navier–Stokes equations using a lattice-gas Boltzmann method, *Europhys. Lett.* 17 (1992) 479.
- [56] Q. Zou, S. Hou, S. Chen, G.D. Doolen, An improved incompressible lattice Boltzmann model for time independent flows, *J. Stat. Phys.* 81 (1995) 35.
- [57] P.L. Bhatnagar, E.P. Gross, M. Krook, A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems, *Phys. Rev.* 94 (1954) 511.
- [58] D. d'Humières, P. Lallemand, Numerical simulations of hydrodynamics with lattice gas automata in two dimensions, *Complex Systems* 1 (1987) 599.
- [59] D.R. Noble, J.G. Georgiadis, R.O. Buckius, Direct Assessment of lattice Boltzmann hydrodynamics and boundary conditions for recirculating flows, *J. Stat. Phys.* 81 (1995) 17.
- [60] R. Mei, W. Shyy, D. Yu, Lattice Boltzmann method for 3-D flows with curved boundary, *J. Comput. Phys.* 161 (2000) 680.
- [61] Z. Guo, C. Zheng, B. Sci, An extrapolation method for boundary conditions in lattice Boltzmann method, *Phys. Fluids* 14 (2002) 2007.
- [62] O. Filippova, D. Hänel, Grid refinement for lattice-BGK models, *J. Comput. Phys.* 147 (1998) 219.
- [63] R.S. Maier, R.S. Bernard, D.W. Grunau, Boundary conditions for the lattice Boltzmann method, *Phys. Fluids* 8 (1996) 1788.
- [64] M. Junk, Z. Yang, Asymptotic analysis of lattice Boltzmann boundary conditions, *J. Stat. Phys.* 121 (2005) 3.
- [65] G. Zhao-Li, Z. Chu-Guang, S. Bao-Chang, Non-equilibrium extrapolation method for velocity and pressure boundary conditions in the lattice Boltzmann method, *Chinese Phys.* 11 (2002) 366.
- [66] Q. Zou, X. He, On pressure and velocity boundary conditions for the lattice Boltzmann BGK model, *Phys. Fluids* 9 (1997) 1591.
- [67] F.M. White, *Viscous Fluid Flow*, McGraw-Hill, New York, 2001, p. 123.
- [68] O. Filippova, B. Schwade, D. Hänel, Multiscale lattice Boltzmann schemes for low Mach number flow, *Phil. Trans. Roy. Soc. London A* 360 (2002) 467.
- [69] N. Thürey, U. Rude, Optimized free surface fluids on adaptive grids with the lattice Boltzmann method, in: SIGGRAPH 2005 Poster.
- [70] J. Tölke, S. Freudinger, M. Krafczyk, An adaptive scheme using hierarchical grids for lattice Boltzmann multi-phase flow simulations, *Comput. Fluids* 35 (2006) 820.
- [71] G.M. Morton, A Computer Oriented Geodetic Data Base and a New Technique in File Sequencing, IBM Ltd, Ottawa, Ontario, 1966.
- [72] D. Kandhai, A. Koponen, A.G. Hoekstra, M. Kataja, J. Timonen, P.M.A. Sloot, Lattice-Boltzmann hydrodynamics on parallel systems, *Comput. Phys. Comm.* 111 (1998) 14.
- [73] A. George, J.W.-H. Liu, *Computer Solution of Large Sparse Positive Definite Systems*, Prentice-Hall, Englewood Cliffs, 1981.
- [74] T. Goehring, Y. Saad, Heuristic algorithms for automatic graph partitioning, Technical Report UMSI- 94-29, University of Minnesota Supercomputing Institute, 1994.
- [75] P. Ciarlet Jr., F. Lamour, On the validity of a front-oriented approach to partitioning large sparse graphs with a connectivity constraint, Technical Report 94-37, Computer Science Department, UCLA, Los Angeles, CA, 1994.
- [76] X. Wu, V. Taylor, S. Garrick, D. Yu, J. Richard, Performance analysis, modeling and prediction of a parallel multiblock lattice Boltzmann application using prophesy system, in: IEEE International Conference on Cluster Computing, Barcelona, Spain, 2006, p. 25.
- [77] S. Dong, G.E. Karniadakis, N.T. Karonis, Cross-site computations on the TeraGrid, *Comput. Sci. Engrg.* 7 (2005) 14.
- [78] B.M. Boghosian, P.V. Coveney, S. Dong, L.I. Finn, S. Jha, G.E. Karniadakis, N.T. Karonis, Nektar, SPICE and Vortronics: Using federated grids for large scale scientific applications, *Cluster Comput.* 10 (2007) 351.
- [79] S. Jha, P.V. Coveney, M. Harvey, SPICE: Simulated Pore Interactive Computing Environment, in: Supercomputing, 2005, Proceedings of the ACM/IEEE SC 2005 Conference.
- [80] B.M. Boghosian, L.I. Finn, P.V. Coveney, Moving the data to the computation: multi-site distributed parallel computation, <http://www.realitygrid.org/publications.shtml>, 2006.
- [81] <http://www3.niu.edu/mpi/>.
- [82] <http://www-unix.mcs.anl.gov/toonen/mpig/>.
- [83] <http://www.globus.org/>.
- [84] J. Chin, J. Harting, S. Jha, P.V. Coveney, A.R. Porter, S.M. Pickles, Steering in computational science: mesoscale modelling and simulation, *Contemp. Phys.* 44 (2003) 417.
- [85] J.M. Tendler, S. Dodson, S. Fields, H. Le, B. Sinharoy, IBM, POWER4 System Microarchitecture, IBM Server Group, <http://www.hpcx.ac.uk/>, 2001.
- [86] <http://www.teragrid.org/userinfo/hardware/resources.php>.
- [87] <http://www.hector.ac.uk/>.
- [88] <http://www.loni.org/systems/>.