

Heterogeneous System-on-Chip based Lattice-Boltzmann Visual Simulation System

Xiaojun Zhai, *Member, IEEE*, Minsi Chen, Sahar Soheilian Esfahani, Abbas Amira, *Senior Member, IEEE*, Faycal Bensaali, *Senior Member, IEEE*, Julien Abinahed, Sarada Dakua, Robin A. Richardson, and Peter V. Coveney

Abstract—Cerebral aneurysm is a cerebrovascular disorder caused by a weakness in the wall of an artery or vein, that causes a localised dilation or ballooning of the blood vessel. It is life-threatening, hence an early and accurate diagnosis would be a great aid to medical professionals in making the correct choice of treatment. HemeLB is a massively parallel lattice-Boltzmann simulation software which is designed to provide the radiologist with estimates of flow rates, pressures and shear stresses throughout the relevant vascular structures, intended to eventually permit greater precision in the choice of therapeutic intervention. However, in order to allow surgeries and doctors to view and visualise the results in real-time at medical environments, a cost-efficient, practical platform is needed. In this paper, we have developed and evaluated a version of HemeLB on various heterogeneous system-on-chip platforms, allowing users to run HemeLB on a low cost embedded platform and to visualise the simulation results in real-time. A comprehensive evaluation of implementation on the Zynq SoC and Jetson TX1 embedded graphic processing unit platforms are reported. The achieved results show that the proposed Jetson TX1 implementation outperforms the Zynq implementation by a factor of 19 in terms of site updates per second.

Index Terms— Zynq; GPU; HemeLB; Lattice Boltzmann; Visualisation; Cerebral Aneurysm

I. INTRODUCTION

MEDICINE and Physiology are being revolutionised through innovations made possible via the growth and application of information technology. The new technology allows the development of increasingly detailed computational models of the biological processes that sustain life, using data from both animals and humans. These models, in conjunction with the relevant experimental data, are helping researchers to gain insights into the physiology and pathology of a multitude of different biological systems, in many cases beyond what is possible through purely observational methods [1], [2].

X. Zhai is with the School of Computer Science and Electronic Engineering, University of Essex, Colchester, UK. e-mail: (xzhai@essex.ac.uk).

M. Chen is with Department of Computer Science, University of Huddersfield, Huddersfield, UK. e-mail:(M.Chen@hud.ac.uk).

A. Amira is with Institute of Artificial Intelligence, De Montfort University, Leicester, UK. e-mail:(abbes.amira@dmu.ac.uk)

S. S. Esfahani, and F. Bensaali are with College of Engineering, Qatar University, Doha, Qatar. e-mail:(sahar.esfahani@qu.edu.qa; f.bensaali@qu.edu.qa).

J. Abinahed and S. Dakua are with Department of Surgery, Hamad Medical Corporation, Doha, Qatar. e-mail:(jabinahed@hamad.qa; sdakua@hamad.qa).

R. Richardson and P. Coveney are with Centre for Computational Science, University College London, London, UK. e-mail:(robin.richardson@ucl.ac.uk; p.v.coveney@ucl.ac.uk).

Manuscript received February 15, 2019; revised September 12, 2019.

Cerebrovascular disorders, such as the rupture of cerebral aneurysms during subarachnoid haemorrhages in Figure 1, represent one of the most prevalent and devastating diseases of adults worldwide. There are noticeable ethnic and geographical differences in its incidence, prevalence, and outcomes. Although the incidence of such disorders tends to decrease in some developed countries, like the UK, where the total burden has been £0.5 billion annually [3], it is increasing in other countries and is assuming epidemic proportions in developing countries [4], [5].




Fig. 1: Cerebral aneurysms [6].

Effective surgical treatments of intracranial aneurysms (ICAs) often make use of endovascular approaches, which are sometimes more effective and maintain low operative risk with short lengths of stay in hospital. These approaches are focused on using intra-aneurysmal coils, and may fail in some cases due to incomplete occlusion of the defect. One solution is to use coils in combination with stents that reorganise the flow of blood into and around the aneurysm. A key factor in the successful deployment of such approach is the skill of the interventional radiologist in identifying the vascular geometry and estimating the consequent fluid flows from inspection of 2D and/or 3D images, in order to propose the best treatment for the patient. Currently, there are few methods to measure these flows and related pressures intraoperatively and the treatment is often based on the experience and intuition of the radiologist.

Computational haemodynamics techniques are widely used to estimate local fluid flow. The computational approach must not only model the native flows in the ICA, but also the perturbations introduced by insertion of flow-diverting stents,

as well as the effect on the pattern of clotting within coil-filled aneurysms. Most usefully, simulation input data should be specific to the patient in question, thus taking into account the variability of vascular geometries, vessel wall mechanics and flow-phase specific pulsatile changes in pressures and shear stresses across patients [7], or differences in physiological states for a given patient (such as heart rate or blood pressure) [8].

The use of a fluid mechanics model requires considerable computing resources for the simulation of what can often be large and complex systems, and the simulation software must efficiently deal with the (highly) sparse geometries that are common to vascular networks. Convention techniques for solving fluid flow are commonly based on finite element methods (FEM). Due to the irregularity in vascular morphology, this often implies a non-trivial task of generating meshes from complex vascular networks, and prescribing additional boundary conditions for resolving fluid-structure interaction.

To this end, we adopted a lattice-Boltzmann (LB) solver for simulating haemodynamics and fluid-vessel interaction. Previous studies showed a comparable accuracy between LB and FEM [9], [10]. In this study, we used HemeLB, a massively parallel LB solver optimized for sparse and complex systems on large supercomputing resources [8], [11], [12]. It has been designed with the ultimate intention of allowing doctors to investigate cerebral blood flow behaviour in the human body and demonstrated its potential in hospital environment [13], [14]. Additionally, HemeLB was designed to meet the requirements of both conventional scientific research software and professional software engineering techniques in an open source way, which allows us to optimise its architecture and re-target to other hardware-accelerated platforms, and evaluate the performance accordingly [15].

HemeLB has been successfully deployed on open academic high-performance computing (HPC) platforms such as HEC-ToR, ARCHER, SuperMUC and Blue Waters [8], [12], [13], among others. In order to allow such hemodynamic simulations to be run as part of routine clinical practice, the workload will need to be deployed on a dedicated computational infrastructure. This is necessary to reduce the dependence of the final computational workflow on a distributed environment not fully under the control of the clinician.

A number of novel architectures: multicore, general purpose graphical progressing units (GPGPUs) and field programmable gate arrays (FPGAs) have been developed recently, which have demonstrated great potential for accelerating computationally intensive tasks such as those found in machine learning and artificial intelligence applications [16] [17]. In this paper, a solution for designing and implementing HemeLB on cost efficient embedded platforms such as to allow visualisation of the simulation results in real-time and the execution of the code in the local environment of a hospital is presented, which is likely necessary for reasons of patient data protection. We first report the evaluation results of our HemeLB simulation software implementation on a cost and energy efficient multi-processor system-on-chip (MPSoC) based Zynq SoC platform. We then present the results for the adaptation of the original HemeLB simulation software to a GPU implementation, which uses a

Jetson TX1 embedded GPU platform. Additionally, a range of comprehensive tests using real patient input data have been carried out with this implementation, and the corresponding performance is reported. The results show that this platform could be the solution to the problem of in-hospital execution and real-time visualisation, offering substantial processing performance on a number of local platforms.

The remaining sections of this paper are organised as follows. In Section 2, HemeLB model is described. Section 3 presents the implementations of HemeLB. Experimental results and performance analysis are reported in section 4 followed by the conclusions in Section 5.

II. HEME LB MODEL

HemeLB uses the LB method to simulate the (continuum) fluid flow, modelled by a grid of particle density distribution functions representing ensembles of particles performing local propagation and collision processes over a discrete lattice mesh [11]. In the following system we use a three dimensional lattice with 19 discrete speeds, called the D3Q19 lattice shown in Figure 2.




Fig. 2: Lattice node of D3Q19 model [18].

Figure 2 illustrates one fluid site and the finite number of discrete lattice vectors, e_i , via which fluid may travel to neighbouring fluid sites in a D3Q19 model. The distribution function is assumed, in the simplest LB implementation, to evolve towards its local equilibrium value according to a single relaxation parameter τ .

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

where f is the distribution function of the particles, and Ω is the Bhatnagar–Gross–Krook (BGK) collision operator. The lattice BGK equation is then formulated as:

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

In (2), the local equilibrium distribution function f_i^{eq} is given as:

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

where τ is the relaxation time towards equilibrium for collisions which is calculated separately from streaming. w_i is a weight coefficient, c_s is the speed of sound, e_i is the particle's velocity in the direction i and the hydrodynamic density p

and macroscopic velocity u are determined by the distribution functions based on

$$p = \sum_{i=0}^{18} f_i = \sum_{i=0}^{18} f_i^{(eq)} \quad (4)$$

$$\mathbf{u} = \sum_{i=0}^{18} \mathbf{e}_i f_i = \sum_{i=0}^{18} \mathbf{e}_i f_i^{(eq)} \quad (5)$$

The distribution functions are propagated along the lattice velocity e_i to the adjacent sites. More specifically, the equilibrium distribution function f_i^{eq} moves from the site at position (x, y, z) to the site at position $(x, y, z) + e_i$. The microscopic velocity in lattice nodes is given as:

$$\vec{e}_i = \begin{cases} (0, 0, 0) & i = 0 \\ (\pm 1, 0, 0), (0, \pm 1, 0), (0, 0, \pm 1) & i = 1, 2, 3, 4, 5, 6 \\ (\pm 1, \pm 1, 0), (\pm 1, 0, \pm 1), (0, \pm 1, \pm 1) & i = 7, 8, 9, \dots, 18 \end{cases} \quad (6)$$

The streaming calculation updates the particle distributions according to the 19 directions velocity f_i . Consequently, the density and velocity are calculated using equations (4) and (5) from f_i . In addition, the equilibrium distribution and the distribution function in the collision step are also calculated. Finally, the streaming and collision steps are repeated.

HemeLB is implemented with various boundary conditions (BC), for example, velocity inlet BC: Ladd iolets [19], Bouzidi-Firdaouss-Lallemand (BFL): interpolated wall collision BC [20], and pressure iolets: mixed Dirichlet-Neumann boundary condition [21].

A. Computational core of lattice-Boltzmann

A two-level data structure is used to represent a domain with a single resolution when doing early load decomposition, which includes a coarser and finer grid layers. The main benefits of this approach are that no data exchange takes place between the coarser grid and finer one, and the structure saves memory with respect to the full matrix representation. The LB algorithm has two main computation procedures: 1) collision phase; and 2) streaming phase, where the collision phase is used to compute the local velocity parameters in n directions (i.e. $n = 19$), and the streaming phase is used to update the velocities with the adjacent sites. Compared with the streaming phase, the collision phase contains the main arithmetic computations of the lattice-Boltzmann approach. This depends on the configuration of the model: for example, if the D3Q19 model is used, each site contains a N^3 cubic lattice with the 3D data stored in 1D arrays, and a lattice grid of integers $is_solid[N \times N \times N]$ keeps the updates of the presence of the boundary conditions for each velocity direction i .

The procedure to handle the collision and streaming phases is illustrated in Algorithm 1. The run-time complexity of Collision and Streaming operation is $O(QN)$, where Q is the number of neighbouring fluid sites of a given node, and N is the number of fluid site. Additionally, since there is no data dependency amongst the fluid nodes in this operation, the parallelization of the routine is relatively trivial.

B. Parallelization strategy in HemeLB

HemeLB simulation software uses domain decomposition as the parallelization strategy to ensure that it has fast and high quality of computational simulations. One of the main improvements made in HemeLB is to use two-level data representation, where the graph growing partitioning (GGP) method is only applied to the coarser grid of the two-level data hierarchy [12]. By applying this, the memory overhead has been significantly reduced, and all the fluid lattice sites of a block are assigned to one of the processors, which would reduce the communication overhead for each fluid lattice site of a block. In addition to this, HemeLB also tries to achieve good computational and communication balance in the distributed computing tasks, which ensures that all the parallel tasks are distributed accordingly on different processing cores. HemeLB has also been optimised to reduce inter-rank communication needed wherever possible, which significantly reduces the amount of intra-machine communication in cross-site runs in grid deployments at each time step. In HemeLB, only the distribution function values are communicated, and each processor calculates its interface-dependent identifiers after domain decomposition is achieved. In addition, the two-level representations are represented by 1D arrays, which further reduces the computational cost of accessing an element of an array, and the extra connectivity buffers help the on-processor propagation of particles between fluid lattice sites.

The output data from HemeLB are the effective pressure, velocity and von Mises stress flow values for all (requested) fluid sites in single precision. In comparison to storing all the distribution functions for every lattice site, this approach requires 6 times lower memory consumption. In addition, HemeLB can be configured with a number of check points, which offers flexibility to store the intermediate simulation results to portable binary format files.

In summary, HemeLB reduces redundant operations, increases pattern regularity, simplifies the computational core and optimises intra-machine communications. It also offers a topology-aware two-level domain decomposition to provide a high quality domain decomposition, ensuring a good workload distribution which is necessary for the parallel scalability of a lattice-Boltzmann simulation.

III. IMPLEMENTATIONS OF HEMELB

Multiprocessor system-on-chips (MPSoCs) are widely deployed in high-performance computing and application-specific embedded systems such as gaming and aerospace for real-time response. Moreover, they offer energy efficiency and performance advantages over uni-processor architectures [22]–[24]. Thus, MPSoCs are becoming the computing engines of choice in embedded systems for real-time applications. A dramatic increase in their use is expected in the coming years and there will likely be hundreds of processors on a single chip [23]. In the following section, two novel implementations of HemeLB on the Zynq-7000 and Jetson TX1 development boards are discussed.

Algorithm 1 Collision and Streaming Phases

Input: f_{old} is source buffers of particle distribution function.
Output: f_{new} is destination buffers of particle distribution function

- 1: **procedure** COLLISIONANDSTREAM(N_x, N_y, N_z, d) $\triangleright N_x, N_y$ and N_z are the dimensions of grid; d is the number of vectors.
- 2: $N \leftarrow N_x N_y N_z$ \triangleright The number of fluid nodes.
- 3: $i \leftarrow 0$
- 4: **while** $i < N$ **do**
- 5: $f_i^{eq} \leftarrow w_i(p + \frac{\mathbf{e}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{e}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2})$ \triangleright Calculate local equilibrium dist. func. f_i^{eq}
- 6: **for all** 19 directions **do**
- 7: **if** site is inside the geometry **then** \triangleright Streaming Operator
- 8: $f_{new}(same_direction) \leftarrow f_{old}$
- 9: **else**
- 10: $f_{new}(opposite_direction) \leftarrow f_{old}$
- 11: **end if**
- 12: **end for**
- 13: $i \leftarrow i + 1$
- 14: **end while**
- 15: **end procedure**

A. Architecture of Zynq Implementation

The Zynq-7000 SoC was introduced to combine the software programmability of an ARM-based dual-core processor with the hardware programmability of an FPGA, which has a good potential to achieve a high performance-per-watt as well as the scalability to meet different application requirements [25]. In this work, we use a Digilent Arty Z7-20 Zynq-7000 development board, equipped with a Zynq-7000 all programmable SoC, with 512 MB DDR3 memory and a 16 GB SD card. In addition, there is a 650 MHz dual-core Cortex-A9 processor together with programmable logic equivalent to Artix-7 FPGA [26]. The Arty Z7-20 has HDMI in/out ports, and audio out as well as a number of GPIO ports, directly connected to programmable logic (PL). In addition, there are various interfaces between processing system (PS) and PL for connecting different IP cores [27].

In this work, we adopt the Pynq framework to support our implementation, in which the PL are presented as hardware libraries that are used to support essential I/O and memory access [28]. An Ubuntu 16.04 Linux operating system is deployed on the PS, which is used to support the essential software building package for HemeLB implementation. The PS/PL interfaces of the proposed implementation are shown in Figure 3, where PL and PS are interfaced by various AXI (advanced extensible interface) high performance (HP) and general purpose (GP) ports [16].

On the PS side, the ARM A9 processors load the PYNQ-Z1 v2.1 image from SD card [28], which runs a 32-bit Ubuntu 16.04 Linux operating system. This image provides High-level software implementation productivity for Zynq, which makes it easier to exploit the features of HemeLB system.

B. Architecture of Jetson Tx1 Implementation

Jetson TX1 module is NVIDIA’s latest processor system-on-module for embedded applications, which has a low-cost Tegra X1 chip. Similar to the Zynq system, the Tegra X1

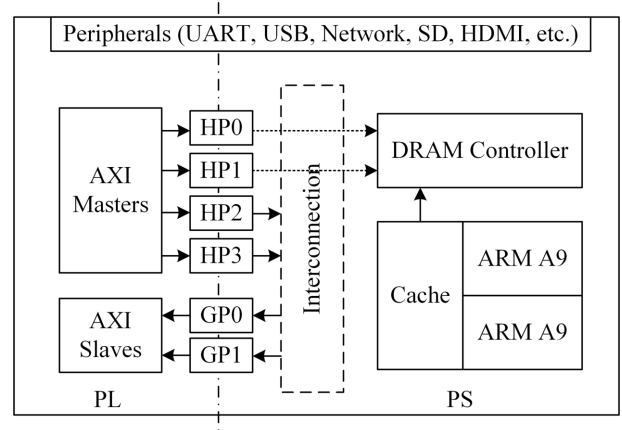


Fig. 3: Zynq Architecture.

CPU subsystem consists of four ARM Cortex-A57 cores [29]. Additionally, the Tegra X1 also has 256 GPU cores based on NVIDIA’s Maxwell architecture. In this work, NVIDIA’s CUDA technology has been used to program the GPU to handle 3D graphics. Figure 4 shows the architecture of the Nvidia Tegra X1 processor.




Fig. 4: Architecture of Nvidia Tegra X1.

In this work, the Jetson Developer Kit is used to evaluate

the performance of HemeLB implementation on the Tegra X1 processor, where JetPack 3.3 for Jetson TX1 has been installed on the kit that runs a 64-bit Ubuntu 16.04 Linux operating system.

C. HemeLB configurations

HemeLB package used in this paper can be obtained from Github website [30]. In addition to this, an *openmpi* package is needed to support the compilation of HemeLB package. Once the entire package is built, we need to set a number of boundary conditions, which are listed in the following table.

TABLE I: HemeLB Boundary Settings

cmake settings	Values
HEMELB_INLET_BOUNDARY	LADDIOLET
HEMELB_WALL_INLET_BOUNDARY	LADDIOLETBFL
HEMELB_USE_VELOCITY_WEIGHTS_FILE	ON
HEMELB_WALL_BOUNDARY	BFL

In addition to the boundary settings specified by the above compilation flags, a list of settings on simulation, geometry and inlets outlets are given in Table II.

TABLE II: HemeLB Simulation, Geometry and Inlets/Outlets Settings

HemeLB settings	Values
Step length	2×10^{-5} s
Voxel size	100×10^{-6}
Inlet velocity	import from a file
Inlet radius	1.9×10^{-3}
Outlet amplitude	0.0 mmHg
Outlet phase	0.0 rad
Outlet period	1.0 s

D. Cerebral aneurysm geometries

As cerebral aneurysms are a very patient specific study, clinicians from Hamad medical corporation (HMC) have carefully chosen images of three subjects of 3D rotational angiography (3DRA), which have been obtained from AXIOM-Artis in HMC. On average, each dataset consisted of 400 slices acquired along the long axes of the subjects. The data include 1 pixel per sample, average slice thickness of 0.3 mm, bit depth of 16, pixel spacing of $0.3 \text{ mm} \times 0.3 \text{ mm}$, and matrix size 512×512 . Figure 5 shows the three STereoLithograph (STL) files used in the experiments.

For all the STLs, the following velocity of blood flow applied in the inlet and used inflow conditions measured in the basilar artery of a patient, the time and velocity is presented in Figure 6.

E. Real-time Visualisation

In order to facilitate real-time interaction and simulation steering, the visualisation of HemeLB is performed *in-situ* directly on dedicated GPUs with compute unified device architecture (CUDA) capability. The communication between

the visualisation client and the LB compute nodes was executed using the existing message passing implementation in HemeLB. The architecture of the proposed visualisation framework is shown in Figure 7, and the pseudo code for the CUDA ray casting kernel function is shown in Algorithm 2.

In Figure 7, the compute nodes mainly handle the computations in HemeLB, and the lattice properties are calculated and cached in each node, then transferred to the master node. This node is also responsible for managing view steering and scheduling lattice data transfer from the compute nodes. The incoming lattice data are shared with CUDA-enabled GPU nodes, and the rendering are performed simultaneously.

In the current prototype, we use the direct volume rendering method for visualising the lattice property fields of the lattice. It is worth noting that the number of voxels used for visualisation does not necessarily match the size of the LB lattice. In common usage for real-time monitoring, the visualisation volume can be 50% or 25% of the simulation resolution. This brought in an additional benefit of reducing the data transfer payload when requesting lattice data from the LB compute nodes.

The internal architecture of the visualisation client is a three-tier system as depicted in Figure 8. The frontend is an OpenGL application. The middle tier is the host layer where application data including cached lattice properties and steering parameters are stored. The top tier is the rendering layer which stores all the voxels for direct volume rendering. The CUDA cores in the top tier execute the volume rendering kernel based on the ray marching algorithm [31].

A number of memory transfers in this architecture are time-critical for real-time visualisation. These include the passing of lattice properties (e.g. velocity, density and pressure) to the visualisation volume and a number of steering simulation. We deployed two level memory access optimisation that can increase throughput between the host and the GPU. Firstly, the visualisation volume was stored in a 3D texture unit which allowed the fast sampling of voxel values by the rendering kernel. It also allowed us to exploit the automatic trilinear interpolation when handling with the spatial resolution disparity between visualisation and simulation.

Secondly, we used the GPU constant memory buffers for storing the viewing parameters and the look-up table for transfer function. This enabled direct memory access (DMA) to be performed thus eliminating the need for additional bandwidth for memory copies. Furthermore, from a GPU thread's perspective, the efficiency of accessing constant memory is comparable to reading from its register.

The current implementation of the visualisation client was tested on the NVIDIA Jetson TX1 development kit. Figure 9 shows the pressure visualisation results of the s10, s11 and s19 data sets.

The current steerable parameters for visualisation include model rotation, zooming, and adjusting the scaling and offset of the transfer functions.

IV. PERFORMANCE ANALYSIS

HemeLB has been benchmarked using simulation domains based on three distinct geometries from Figure 5. An overview

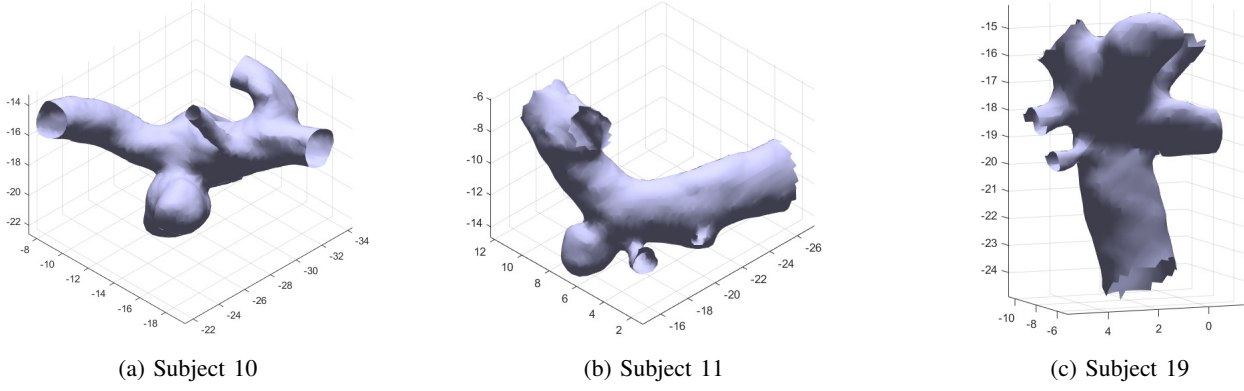


Fig. 5: The surface meshes used in our experiments.

Algorithm 2 Pseudo code for the CUDA ray casting kernel function.

Input: *VolumeData*, *TransferFunction*, *ImageWidth*, *ImageHeight*

Output: *OutputColour*

```

1: procedure VOLUMERENDERKERNEL( $x, y, S, \delta s, \rho, \mathbf{M}$ )  $\triangleright x, y$  are pixel position;  $S$  is ray marching step;  $\delta s$  is ray step
   size;  $\rho$  is voxel density;  $\mathbf{M}$  is the view matrix
2:    $OutputColour \leftarrow 0$ 
3:    $u \leftarrow x/ImageWidth * 2 - 1$   $\triangleright (u, v)$  are normalised screen coordinates
4:    $v \leftarrow y/ImageHeight * 2 - 1$ 
5:    $\mathbf{o} \leftarrow (0, 0, 0)$   $\triangleright$  Ray origin in canonical view space
6:    $\mathbf{d} \leftarrow (u, v, -2)$ 
7:    $\mathbf{d}' \leftarrow \mathbf{M}^{-1}\mathbf{d}$   $\triangleright M$  is the view matrix
8:    $Ray \leftarrow (\mathbf{o}, \mathbf{d}')$   $\triangleright$  Ray defined in view space
9:    $t \leftarrow INTERSECTVOLUME(VolumeData, Ray)$   $\triangleright$  Ray interests with the nearest voxel in Volume
10:  if  $0 \leq t < MAX$  then  $\triangleright$  Position of the nearest voxel
11:     $\mathbf{p} \leftarrow \mathbf{o} + \mathbf{d}'t$   $\triangleright$  Ray cast step vector
12:     $\delta \mathbf{d} \leftarrow \mathbf{d}'\delta s$ 
13:     $s \leftarrow 0$ 
14:    while  $s < S$  do  $\triangleright$  March along the ray for  $S$  steps
15:       $Value \leftarrow SAMPLEVOLUME(p)$   $\triangleright$  Sample the value at  $\mathbf{p}$ 
16:       $Colour \leftarrow TRANSFERFUNCTION(Value)$   $\triangleright$  Transfer  $Value$  to RGBA
17:       $Colour \leftarrow Colour \cdot \rho$   $\triangleright$  Translucency is determined by  $\rho$ 
18:       $OutputColour \leftarrow ALPHABLEND(OutputColour, Colour)$   $\triangleright$  Integrate voxel colours
19:       $\mathbf{p} \leftarrow \mathbf{p} + \delta \mathbf{d}$   $\triangleright$  March along the ray by  $\delta \mathbf{d}$ 
20:       $s \leftarrow s + 1$ 
21:    end while
22:  end if
23: end procedure

```

of the simulation domain used in the experiments is presented in Table III.

TABLE III: Overview of the Simulation Domains

Name	Number of lattice sites	Number of blocks
Subject 10	186463	2880
Subject 11	151700	2520
Subject 19	81952	728

In Table III, the simulation were performed using a 19-directional LB kernel (D3Q19), the Lattice BKG model [32] with simple bounce-back boundary condition and a fixed physical viscosity of 0.004 Pa.s. In addition, the inlet velocities used in the simulations are from Figure 6. Number of lattice

sites for each simulation subjects are shown in Table III. In the Zynq platform simulation, the timing data is illustrated in Table IV, achieving a maximum performance of 215,751 site updates per second (SUPS) with 2 cores.

TABLE IV: Timing Data for Simulation on Zynq Platform

Name	Total time (s)	Number of steps	SUPS
Subject 10	38,900	45,010	215,751
Subject 11	36,900	52,232	214,732
Subject 19	17,700	45,374	210,084

As it can be seen from Table IV, the timing performances for different subject samples are similar, however, for the large geometry subject, the SUPS is slightly higher than the small




Fig. 6: The peak blood velocity profile used as input to HemeLB




Fig. 7: The architecture of the proposed visualisation framework

geometry subject. This means that the SUPS per core is largely independent of other factors.

We present timing profiles for each subject using Zynq SoC in Figure 10. As it can be seen from the comparisons in Figure 10, the ratios from each time profile are similar, this is because the complexity of algorithm 1 is linear, and the computational resources were constant, therefore, by increasing the lattice sites of the simulation, the proportions of each time profile were the same. The timing consumed for LB computation is slightly less than MPI communication, and both processes consume 97% of the overall processing time in the simulations.

Similar experiments have been conducted on the Jetson TX1 platform for performance evaluations. The timing data is illustrated in Table V, achieving a maximum performance of 4,154,802 site updates per second (SUPS) with 4 cores. In comparison with the SUPS of the Zynq platform, the results from the Jetson TX1 are significantly improved, which is to say, about 19 times better than the Zynq implementation. This is mostly due to the capacities of the processors and memory on Jetson TX1 and PYNQ-Z1 boards, as the former uses a quad-core ARM Cortex-A57 processor at 2 GHz with 4 GB available memory compared to a Dual-core ARM Cortex A9 processor at 650 MHz with only 512 MB available memory. However, in terms of the SUPS results for different subjects, the Jetson TX1 platform has shown very similar

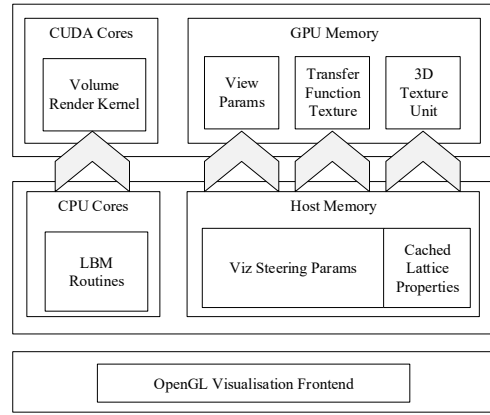


Fig. 8: The internal architecture of visualization client

performance compared to the Zynq platform, as the subjects with larger number of lattice sites achieved slightly higher SUPS compared to other subjects.

TABLE V: Timing Data for Simulation on Jetson TX1 Platform

Name	Total time (s)	Number of steps	SUPS
Subject 10	2,020	45,010	4,154,802
Subject 11	1,930	52,232	4,105,489
Subject 19	907	45,374	4,099,768

We also present timing profiles for each subject using Jetson TX1 in Figure 11. As can be seen from the results in Figure 11, the ratios from each time profile are similar, the timing consumed for LB computation is much higher than message passing interface (MPI) communication, this is different to the Zynq platform, which means that HemeLB works better on the Jetson TX1 platform. Overall, both LB and MPI communication processes consume 96% of the overall processing time in the simulations, which is close to the results from Zynq platform.

A list of memory throughput for each subject are reported in Figure 12. In Figure 12 (a), we have compared the memory throughput of CUDA Host-to-Device (HtoD) for S10, S11, and S19, the peak throughput 176.06 MB/s is achieved in S19, and with an average throughput of 108.35 MB/s. The CUDA Memset (i.e. Fill block of memory in CUDA) for each subject are compared in Figure 12 (b), the throughput achieved in each subject are similar, where the peak throughput is 2.0×10^4 GB/s, and the average throughput is 1.477×10^4 GB/s Figure 12 (c) shows the memory throughput of CUDA Host-to-Array (HoA), the peak throughput is 2.22 GB/s, and average throughput is 1.99 GB/s.

V. CONCLUSION

In this paper, we presented the design and implementation of the HemeLB simulation software on low cost MPSoC embedded platforms. The proposed solutions allow the HemeLB simulation software to be implemented in a local environment rather than a distributed environment, which is not fully under the clinicians control. A set of comprehensive tests using real

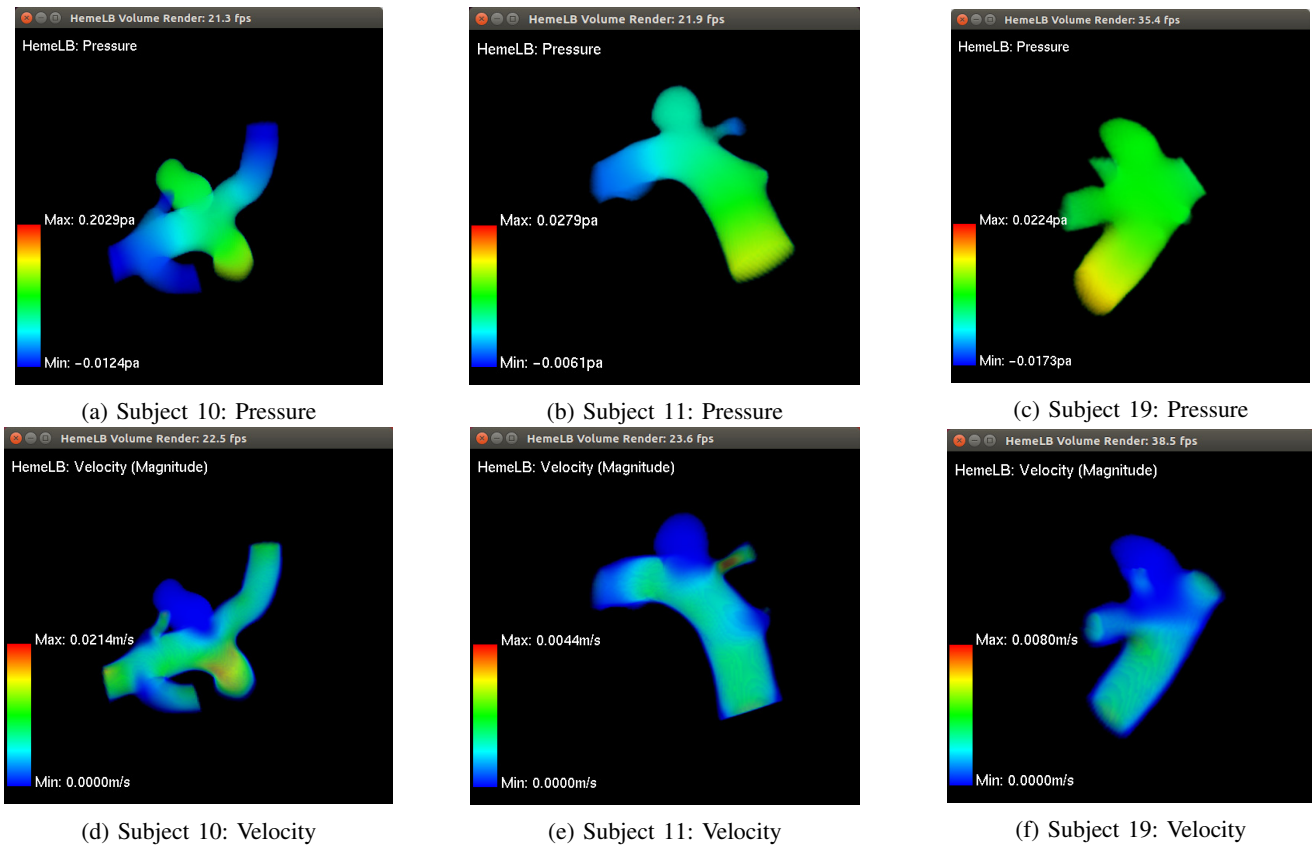


Fig. 9: Visualisation results from the visualisation client. The pressures at each lattice site was shown using an interactively adjustable transfer function.




Fig. 10: Timing profile of each subject using Zynq SoC.

patients data have been carried out on the implementations, and the corresponding evaluation of the performance of the implementation has been also reported. In addition, real-time visualisation of the simulation results are also reported in this paper. According to the evaluation results, Jetson TX1 based SoC is more suitable for implementing HemeLB system than Zynq based SoC, one of the reason is because the CUDA cores and multicore processor in Jetson TX1 offers better hardware infrastructure for accelerating real-time visualisation applications.

In future, we will develop a user friendly graphic user interface (GUI) to allow hospital users to manipulate the operations of the framework and steer the visualisation, to use the current system as a test platform in hospital environments.

ACKNOWLEDGMENT

This paper was made possible by National Priorities Research Program (NPRP) grant No. 5-792-2-328 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

REFERENCES

- [1] S. Okami and N. Kohtake, "Transitional complexity of health information system of systems: Managing by the engineering systems multiple-domain modeling approach," *IEEE Systems Journal*, pp. 1–12, 2018.
- [2] A. A. Abdellatif, A. Mohamed, and C. Chiasserini, "User-centric networks selection with adaptive data compression for smart health," *IEEE Systems Journal*, vol. 12, no. 4, pp. 3618–3628, Dec 2018.




Fig. 11: Timing profile of each subject using Jetson Tx1.




Fig. 12: Memory throughput for each subject using Jetson Tx1.

[3] P. Scarborough, P. Bhatnagar, K. K. Wickramasinghe, S. Allender, C. Foster, and M. Rayner, "The economic burden of ill health due to diet, physical inactivity, smoking, alcohol and obesity in the UK: An update to 2006-07 NHS costs," *Journal of Public Health*, vol. 33, no. 4, pp. 527–535, 2011.

[4] N. Benkirane, A. Bennis, S. Bellakhdar, R. Habbal, S. Aidi, M. El Alaoui Faris, H. El Otmani, B. El Moutawakil, M. Rafai, and I. Slassi, "Stroke risk factors in a Moroccan population: a multicentric prospective study," *Journal of the Neurological Sciences*, vol. 357, p. e367, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0022510X15018109>

[5] D. Lin, Y. Tang, and A. V. Vasilakos, "User-priority-based power control in d2d networks for mobile health," *IEEE Systems Journal*, vol. 12, no. 4, pp. 3142–3150, Dec 2018.

[6] Steps to Health, "Brain Aneurysms: What Are They and How to Prevent Them." [Online]. Available: <https://steptohealth.com/brain-aneurysms-prevent/>

[7] B. M. Jobst, R. Hindriks, H. Laufs, E. Tagliazucchi, G. Hahn, A. Ponce-Alvarez, A. B. A. Stevner, M. L. Kringelbach, and G. Deco, "Increased stability and breakdown of brain effective connectivity during slow-wave sleep: mechanistic insights from whole-brain computational modelling," *Scientific reports*, vol. 7, no. 1, p. 4634, 2017.

[8] A. Patronis, R. A. Richardson, S. Schmieschek, B. J. Wylie, R. W. Nash, and P. V. Coveney, "Modelling patient-specific magnetic drug targeting within the intracranial vasculature," *Frontiers in physiology*, vol. 9, p. 331, 2018.

[9] D. Kandhai, D. J. Vidal, A. G. Hoekstra, H. Hoefsloot, P. Iedema, and P. M. Sloot, "Lattice-Boltzmann and finite element simulations of fluid flow in a SMRX static mixer reactor," *International Journal for Numerical Methods in Fluids*, 1999.

[10] L. Axner, A. G. Hoekstra, A. Jeays, P. Lawford, R. Hose, and P. M. Sloot, "Simulations of time harmonic blood flow in the Mesenteric artery: comparing finite element and lattice Boltzmann methods," *Biomedical engineering online*, 2009.

[11] M. D. Mazzeo and P. V. Coveney, "HemeLB: A high performance parallel lattice-Boltzmann code for large scale fluid flow in complex geometries," *Computer Physics Communications*, vol. 178, no. 12, pp. 894–914, 2008.

[12] D. Groen, J. Hetherington, H. B. Carver, R. W. Nash, M. O. Bernabeu, and P. V. Coveney, "Analysing and modelling the performance of the HemeLB lattice-Boltzmann simulation environment," *Journal of Computational Science*, vol. 4, no. 5, pp. 412–422, 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.jocs.2013.03.002>

[13] D. Groen, R. A. Richardson, R. Coy, U. D. Schiller, H. Chandrashekar, F. Robertson, and P. V. Coveney, "Validation of patient-specific cerebral blood flow simulation using transcranial doppler measurements," *Frontiers in Physiology*, vol. 9, p. 721, 2018.

[14] H. Shi, J. Chen, W. Pan, K. Hwang, and Y. Cho, "Collision avoidance for redundant robots in position-based visual servoing," *IEEE Systems Journal*, pp. 1–11, 2018.

[15] UCL, "GitHub - UCL/hemelb," 2019. [Online]. Available: <https://github.com/UCL/hemelb>

[16] X. Zhai, M. Eslami, E. S. Hussein, M. S. Filali, S. T. Shalaby, A. Amira, F. Bensaali, S. Dakua, J. Abinahed, A. Al-Ansari, and A. Z. Ahmed, "Real-time automated image segmentation technique for cerebral aneurysm on reconfigurable system-on-chip," *Journal of Computational Science*, vol. 27, pp. 35–45, 7 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S187750317313005>

[17] H. Djelouat, X. Zhai, M. A. Disi, A. Amira, and F. Bensaali, "System-on-Chip Solution for Patients Biometric: A Compressive Sensing-Based Approach," *IEEE Sensors Journal*, p. 1, 2018.

[18] X. Zhai, A. Amira, F. Bensaali, A. Al-Shibani, A. Al-Nassr, A. El-Sayed, M. Eslami, S. P. Dakua, and J. Abinahed, "Zynq soc based acceleration of the lattice boltzmann method," *Concurrency and Computation: Practice and Experience*, vol. 0, no. 0, p. e5184, e5184 CPE-18-0066.R1. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.5184>

[19] A. J. Ladd, "Numerical simulations of particulate suspensions via a discretized boltzmann equation. part 1. theoretical foundation," *Journal of fluid mechanics*, vol. 271, pp. 285–309, 1994.

[20] M. Bouzidi, M. Firdaouss, and P. Lallemand, "Momentum transfer of

- a boltzmann-lattice fluid with boundaries,” *Physics of fluids*, vol. 13, no. 11, pp. 3452–3459, 2001.
- [21] R. W. Nash, H. B. Carver, M. O. Bernabeu, J. Hetherington, D. Groen, T. Krüger, and P. V. Coveney, “Choice of boundary condition for lattice-boltzmann simulation of moderate-reynolds-number flow in complex domains,” *Physical Review E*, vol. 89, no. 2, p. 023303, 2014.
- [22] J.-J. Han, M. Lin, D. Zhu, and L. T. Yang, “Contention-Aware Energy Management Scheme for NoC-Based Multicore Real-Time Systems,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 3, pp. 691–701, 3 2015. [Online]. Available: <http://ieeexplore.ieee.org/document/6774493/>
- [23] T. Chantem, R. P. Dick, and X. S. Hu, “Temperature-Aware Scheduling and Assignment for Hard Real-Time Applications on MPSoCs,” in *2008 Design, Automation and Test in Europe*. IEEE, 3 2008, pp. 288–293. [Online]. Available: <http://ieeexplore.ieee.org/document/4484694/>
- [24] X. Zhang, K. Huang, M. Yu, X. Jiang, and X. Yan, “BFCCO: A BPSO-Based Fine-Grained Communication Optimization Method for MPSoC,” *IEEE Access*, vol. 6, pp. 18771–18785, 2018. [Online]. Available: <http://ieeexplore.ieee.org/document/8307752/>
- [25] X. Zhai, A. A. S. Ali, A. Amira, and F. Bensaali, “MLP Neural Network Based Gas Classification System on Zynq SoC,” *IEEE Access*, vol. 4, pp. 8138–8146, 2016. [Online]. Available: <http://ieeexplore.ieee.org/document/7605493/>
- [26] L. H. L. H. Crockett, R. A. Elliot, M. A. Enderwitz, and R. W. E. e. Stewart, *The Zynq book : embedded processing with the ARM Cortex-A9 on the Xilinx Zynq-7000 all programmable SoC*. [Online]. Available: <https://dl.acm.org/citation.cfm?id=2685817>
- [27] Digilent, “Arty Z7 Reference Manual [Reference.Digilentinc].” [Online]. Available: <https://reference.digilentinc.com/reference/programmable-logic/arty-z7/reference-manual>
- [28] Pynq, “PYNQ - Python productivity for Zynq - Home.” [Online]. Available: <http://www.pynq.io/>
- [29] NVIDIA, “NVIDIA [®] Tegra [®] X1 NVIDIA’S New Mobile Superchip,” Tech. Rep., 2015. [Online]. Available: <https://international.download.nvidia.com/pdf/tegra/Tegra-X1-whitepaper-v1.0.pdf>
- [30] UCL, “Hemelb Github repository.” [Online]. Available: <https://github.com/UCL/hemelb>
- [31] K. Zhou, Z. Ren, S. Lin, H. Bao, B. Guo, and H.-Y. Shum, “Real-time smoke rendering using compensated ray marching,” in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3. ACM, 2008, p. 36.
- [32] P. L. Bhatnagar, E. P. Gross, and M. Krook, “A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems,” *Physical Review*, vol. 94, no. 3, pp. 511–525, 5 1954. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.94.511>