Cloud Computing at Scale: Tracking 4.5 Million Heartbeats of 3D Coronary Flow via the Longitudinal Hemodynamic Mapping Framework

CYRUS TANADE, Duke University, USA
AMANDA RANDLES (ADVISOR), Duke University, USA

Tracking hemodynamic responses to treatment and stimuli for long periods is a grand challenge. Moving from established single-heartbeat technology to longitudinal profiles would require continuous data reflecting a patient's evolving state, methods to extend the temporal domain that could be feasibly computed, and high-throughput resources. Although personalized models can accurately measure 3D hemodynamics over single heartbeats, state-of-the-art methods would require centuries of runtime on leadership-class systems to simulate one day of activity. We are establishing the Longitudinal Hemodynamic Mapping Framework (LHMF), which combines patient-specific models, wearables, and cloud computing to enable the first digital twins that capture longitudinal hemodynamic maps (LHMs). We demonstrate validity through comparison with ground truth data for 750 beats. We applied LHMF to generate the first LHM of coronary arteries spanning 4.5 million heartbeats. LHMF relies on an initial fixed set of representative simulations to enable the computationally tractable creation of LHM over heterogeneous systems.

CCS Concepts: • **Computer systems organization → Cloud computing**; • **Applied computing → Health informatics**.

Additional Key Words and Phrases: longitudinal hemodynamic maps, cloud computing, wearable sensors, computational fluid dynamics, cardiovascular disease

1 INTRODUCTION

Personalized 3D blood flow models of coronary arteries are limited to single heartbeat metrics due to their high computational cost. Simulating one day of activity using explicit methods would take centuries on leadership-class systems. To our knowledge, the longest contiguous 3D simulation has only recovered 30 heartbeats [1]. To this end, we propose the Longitudinal Hemodynamic Mapping Framework (LHMF), an approach that combines personalized models, wearable devices, and cloud computing to monitor hemodynamic changes over substantial periods. LHMF provides an additional level of parallelism to our massively parallel fluid solver, HARVEY [2].

Here, we created the first digital twins that capture longitudinal hemodynamic maps (LHM) that cover 4.5 million heartbeats. LHMs capture the full range of behavior longitudinally, unlike conventional hemodynamic maps, which show only one snapshot at a time. We enable LHMs by 1) developing methods to extend the temporal domain that could be tractably simulated, 2) driving digital twins from wearables, 3) validating LHMF with explicit ground truths, and 4) efficiently leveraging high-throughput resources.

2 METHODS

2.1 Deconstructing temporal dependence to enable parallelism in time

We define hemodynamic units (HUs) as the region containing the preceding cardiac cycles used for the convergence and heartbeat of interest (**Figure 1A**). Solutions are considered to temporally converge with a 1% discrepancy threshold

---

Authors' addresses: Cyrus Tanade, cyrus.tanade@duke.edu, Duke University, 534 Research Dr., Durham, North Carolina, USA, 27705; Amanda Randles (ADVISOR), amanda.randles@duke.edu, Duke University, 534 Research Dr., Durham, North Carolina, USA, 27705.

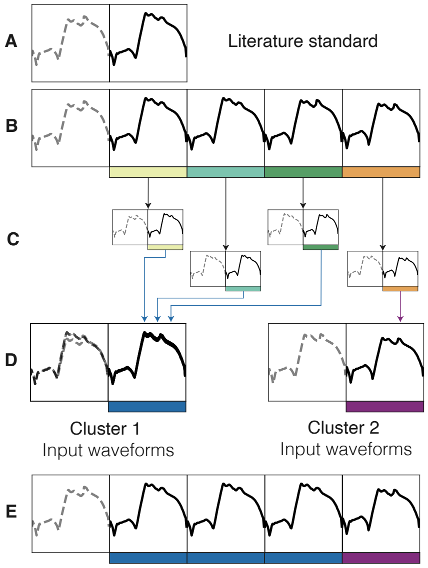

Fig. 1. **Longitudinal Hemodynamic Mapping Framework (LHMF).** Solid lines indicate the heartbeat of interest used for analysis and interrupted gray lines delineate preflow cardiac cycles for temporal convergence. (A) Literature standard for simulating transient flow. (B) Explicit approach for simulating multiple heartbeats. (C) Sequences of heartbeats are deconstructed into hemodynamic units (HUs), which are simulated in parallel via cloud instances. (D) Capturing representative sets of hemodynamically equivalent HUs can minimize redundancy. (E) Clustered HUs are deployed on cloud instances and used to reassemble a longitudinal sequence.

[3]. Longitudinal simulations could be naively performed explicitly, where flow across a long series of heartbeats is simulated contiguously (**Figure 1B**), but is computationally intractable.

We established the LHMF where each heartbeat is mapped to an HU and simulated in parallel (**Figure 1C**). To validate the LHMF, we (1) evaluated temporal convergence across varying activity states to ensure that the solutions converged for all HU and (2) performed a 750 heartbeat simulation using explicit methods to provide a robust ground truth to validate the LHMF. We focus on computing the resting pressure gradient (PG), a metric correlated with atherosclerosis [4], blood velocity (V), and wall shear stress (WSS), the frictional force of blood on the inner walls of blood vessels.

Although LHMF greatly increases the number of heartbeats that can be captured, it is still limited when simulating every heartbeat in a prohibitively long sequence. To overcome this challenge, we established LHMF_C where we clustered hemodynamically similar HUs to minimize performing redundant simulations (**Figure 1D**). Representative HUs could be simulated as a one-time, fixed cost to reconstruct a long sequence of heartbeats (**Figure** 1E).

2.2 Driving 3D coronary digital twins from continuous wearable data

We capture varying activity states longitudinally by augmenting boundary conditions (**Figure 2**). We incorporated continuous physiological data to modify the velocity waveforms on demand to reflect the evolving physical states. LHMF provides an additional level of parallelism to massively parallel 3D fluid solvers.

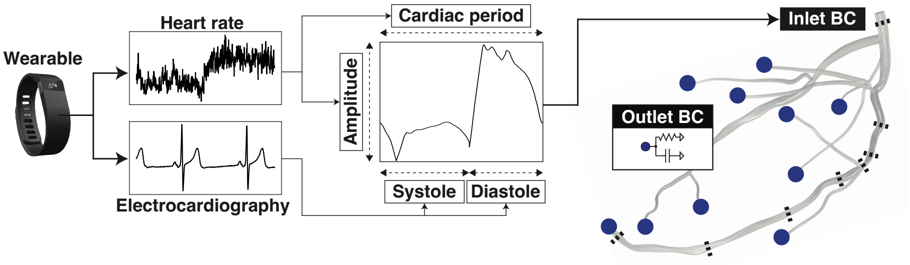

Fig. 2. **Wearable-driven blood flow models.** Instantaneous heart rate and electrocardiography (ECG) data are used to scope out a set of patient-specific 3D computational fluid dynamics simulations.

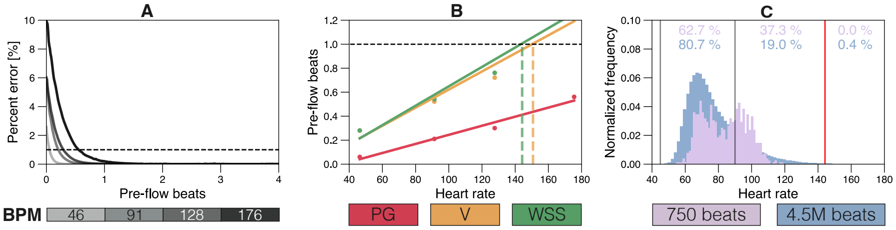

Fig. 3. **Ensuring temporal convergence for all activity states.** (A) Per-time step percentage discrepancy curves for resting pressure gradient. (B) Linear regressions yield R^2 values exceeding 0.97 for all three hemodynamic variables. Vertical interrupted lines delineate the heart rate at which two preflow cycles are needed. (C) Histograms displaying the distribution of heart rates for the 750-beat and 4.5 million-beat cases. The subplot is divided into three bins showing prevalence.

2.3 Deploying simulations on cloud and traditional systems

`hpc6a.48xlarge` instances with 38,400 AMD EPYC 7003 cores (AWS CPUs) were used to generate LHMs. We also used `c7g.16xlarge` (ARM Graviton3 CPUs), `p3dn.24xlarge` (Tesla V100 GPUs), OLCF Summit, TACC Stampede2, and the Duke Compute Cluster.

3 RESULTS

3.1 Few cardiac cycles are required for temporal convergence in varying activity states

Figure 3A demonstrates that error curves shifted to the right with increasing heart rate. When relating the error to the heart rate, the linear regressions indicated near-perfect correlations ($R^2 > 0.97$) (**Figure 3B**). PG required one preflow cycle, while V and WSS required two preflow cycles after around 140 *bpm*.

3.2 LHMF agrees with ground truth explicit methods evaluated with 750 heartbeats

LHMF was validated with the ground truth 750-beat data. All hemodynamic metrics had negligible errors (**Figure 4**).

3.3 LHMF$_C$ drastically improves computational tractability across heterogeneous systems

Using a problem size of 1.5 months, we projected wallclock times in heterogeneous systems (**Figure 5**). The explicit approach required over 10^5 days of compute time. In general, LHMF resulted in over a two orders of magnitude reduction in wallclock time. LHMF$_C$ further decreased runtime to approximately 43 hours between systems.

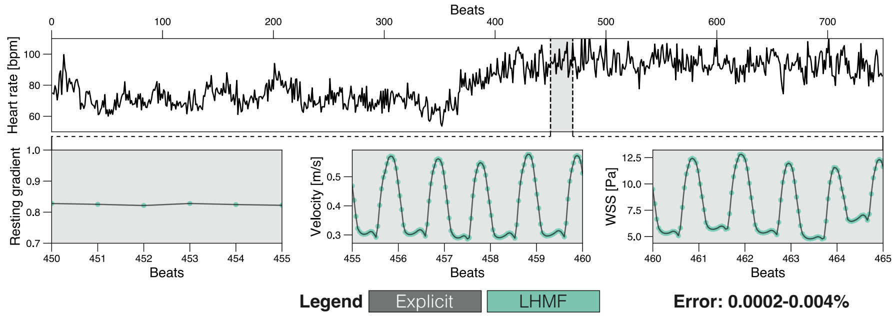

Fig. 4. **LHMF recapitulates ground truth simulations with over 99.995% accuracy.** (Top) Continuous wearable heart rate data over 750 beats. (Bottom) Insets between beats 450-465 compare resting pressure gradient, velocity, and wall shear stress (WSS) using the explicit approach vs. LHMF.

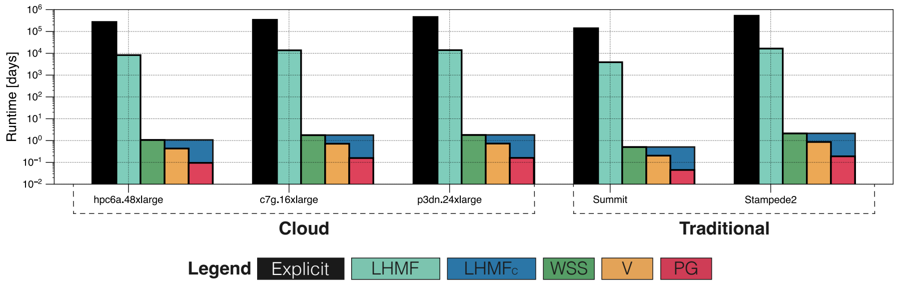

Fig. 5. **LHMF and LHMF$_C$ deployed on heterogeneous resources.** Performance in terms of wallclock runtime across multiple types of cloud instances and traditional clusters. Each grouped bars represent a different system, including `hpc6a.48xlarge` (x86 CPUs), `c7g.16xlarge` (ARM CPUs), and `p3dn.24xlarge` (GPUs) instances. On the traditional cluster side, we present results on OLCF Summit and TACC Stampede2. Results for LHMF$_C$ are subdivided by hemodynamic metric.

3.4 Application of LHMF$_C$ to generate LHMs that span over 4.5 million heartbeats

We focused on generating the first LHM for WSS since low WSS ($< 1\ Pa$) has been linked to the progression of atherosclerosis [5]. We generated spatial maps of WSS for single heartbeats in the rest and exercise state. For comparison, we created a spatial map averaged over 4.5 million heartbeats (or 1.5 months of activity) (**Figure 6A**). The WSS computed from the LHM was statistically different from the single rest and exercise heartbeat results ($p < 0.001$), demonstrating that single heartbeat maps cannot capture longitudinal hemodynamic behavior. In addition, we created a temporal map that shows the percentage of time spent with low WSS, providing a novel way to identify local regions with potential risks for disease development (**Figure 6B**).

4 CONCLUSION AND FUTURE WORK

This study has established frameworks for capture of longitudinal hemodynamics, with LHMF$_C$ producing LHMs on the order of months for the first time. We demonstrated that using a finite set of representative HUs is sufficient to accurately reconstruct the solution of an explicit 3D simulation. In future work, further exploration of different physiological boundary conditions would be needed. We expect LHMs to advance understanding of cardiovascular disease, creating a step-change improvement in the number of heartbeats that can be accurately simulated.

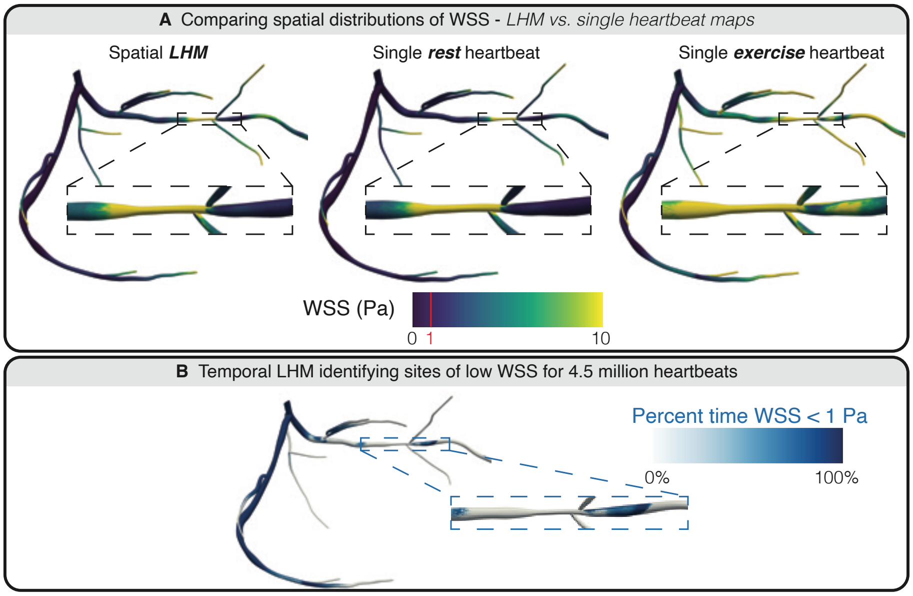

Fig. 6. **Application of LHMF$_C$ for capturing 1.5 months of daily activity (or 4.5 million heartbeats).** (A) Comparing spatial distributions of WSS between an average spatial map of WSS for 4.5 million heartbeats, rest state single heartbeat WSS map, and exercise state single heartbeat WSS map. (B) Temporal LHM of WSS. The fraction of time spent with WSS $< 1\ Pa$ over 1.5 months of activity is visualized in the temporal LHM.

ACKNOWLEDGMENTS

Computing support came from Amazon Web Services (AWS) Cloud Credits for Research, Duke Compute Cluster, and Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) which is supported by NSF grant number ACI-1548562. An award of compute time was provided by the INCITE program. This research used resources from the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. Research reported in this work was supported by the National Institute on Aging of the NIH under Award Number DP1AG082343, NSF GRFP under Grant No. DGE 164486, and the Coulter Foundation. The content does not necessarily represent the official views of the NIH or NSF.

REFERENCES

[1] Noelia Grande Gutiérrez, Mark Alber, Andrew M Kahn, Jane C Burns, Mathew Mathew, Brian W McCrindle, and Alison L Marsden. Computational modeling of blood component transport related to coronary artery thrombosis in kawasaki disease. *PLoS computational biology*, 17(9):e1009331, 2021.

[2] Amanda Peters Randles, Vivek Kale, Jeff Hammond, William Gropp, and Efthimios Kaxiras. Performance analysis of the lattice boltzmann model beyond navier-stokes. In *2013 IEEE 27th International Symposium on Parallel and Distributed Processing*, pages 1063–1074. IEEE, 2013.

[3] Martin R Pfaller, Jonathan Pham, Nathan M Wilson, David W Parker, and Alison L Marsden. On the periodicity of cardiovascular fluid dynamics simulations. *Annals of biomedical engineering*, 49(12):3574–3592, 2021.

[4] Barry Hennigan, Keith G Oldroyd, Colin Berry, Nils Johnson, John McClure, Peter McCartney, Margaret B McEntegart, Hany Eteiba, Mark C Petrie, Paul Rocchiccioli, et al. Discordance between resting and hyperemic indices of coronary stenosis severity: the verify 2 study (a comparative study of resting coronary pressure gradient, instantaneous wave-free ratio and fractional flow reserve in an unselected population referred for invasive angiography). *Circulation: Cardiovascular Interventions*, 9(11):e004016, 2016.

[5] Habib Samady, Parham Eshtehardi, Michael C McDaniel, Jin Suo, Saurabh S Dhawan, Charles Maynard, Lucas H Timmins, Arshed A Quyyumi, and Don P Giddens. Coronary artery wall shear stress is associated with progression and transformation of atherosclerotic plaque and arterial remodeling in patients with coronary artery disease. *Circulation*, 124(7):779–788, 2011.