An Entropy-Bounded Thermodynamic Approach for Real-Time Maternal Hemorrhage Monitoring

Steven D. Harris
Francesca Bonetta-Misteli
Lleyton Martin
Christopher D. Gill
Roger D. Chamberlain
Christine M. O'Brien

Steven D. Harris, Francesca Bonetta-Misteli, Lleyton Martin, Christopher D. Gill, Roger D. Chamberlain, and Christine M. O'Brien, "An Entropy-Bounded Thermodynamic Approach for Real-Time Maternal Hemorrhage Monitoring," in *Proc. of IEEE High-Performance Extreme Computing Conference* (HPEC), September 2026.

Dept. of Computer Science and Engineering
Dept. of Biomedical Engineering
Washington University in St. Louis

An Entropy-Bounded Thermodynamic Approach for Real-Time Maternal Hemorrhage Monitoring

Steven D. Harris*, Francesca Bonetta-Misteli†, Lleyton Martin†,
Christopher D. Gill*, Roger D. Chamberlain*, and Christine M. O'Brien†‡
*Dept. of Computer Science and Engineering
†Dept. of Biomedical Engineering
‡Dept. of Obstetrics and Gynecology
Washington Univ. in St. Louis, St. Louis, Missouri 63130
Email: {sharris22,fbonetta-misteli,l.v.martin,cdgill,roger,c.obrien}@wustl.edu

***Abstract*—To combat postpartum hemorrhage, the cause of 30% of maternal deaths, this research introduces an Entropy-Bounded Performance Model (EBPM) to optimize laser speckle contrast imaging on wearable hardware. Derived from the Boltzmann uniqueness theorem, the EBPM resolves the perceived dichotomy between throughput and latency by establishing that frame rates are fundamentally limited by the Shannon entropy of processing delays. Although a Raspberry Pi 5 prototype achieves 80 FPS, it currently fails clinical safety targets due to moderate jitter and "entropy collapse" events at 85% CPU utilization. The study concludes that deterministic scheduling and multi-core optimization are mandatory to minimize processing entropy, thereby enabling the deployment of low-cost wearables for pre-symptomatic hemorrhage detection in both high- and low-resource settings.**

***Index Terms*—laser speckle contrast imaging, Boltzmann uniqueness theorem, entropy-bound performance model**

I. Introduction

POSTPARTUM hemorrhage (PPH) constitutes a critical global health challenge, representing approximately 30% of maternal deaths worldwide and over 10% in high-resource regions [1], [2]. Current clinical detection methods rely on visual blood loss estimation and vital sign monitoring, which can often be slow and inaccurate due to physiological compensatory mechanisms that maintain core perfusion while reducing peripheral blood flow [3]–[5]. Laser speckle contrast imaging (LSCI) has emerged as a promising non-invasive technique for continuous perfusion monitoring, exploiting the inverse relationship between speckle pattern blur and red blood cell velocity [6].

This work analyzes a prototype LSCI platform implemented on Raspberry Pi 5 hardware with camera module and 785 nm laser diode for peripheral blood flow monitoring in postpartum mothers. The eventual goal is to deploy a wearable system for continuous monitoring. Prior to clinical deployment, a number of hardware/software tradeoffs need to be assessed, including, what fraction of the signal processing should be done on the wearable vs. delivering raw unprocessed signals to a remote system. The present prototype is being used to evaluate the viability of performing the majority of the signal processing on the wearable system itself.

This work supported by NIH grant R00HD103954 and a gift from BECS Technology, Inc.

While accessible low-cost sensing enables deployment in both high- and low-resource settings, the embedded implementation presents significant performance modeling challenges that existing frameworks inadequately address [7], [8]. We identify fundamental contradictions in prevailing computer engineering performance paradigms as applied to real-time biomedical sensing and resolve them through first-principles thermodynamic analysis.

Our contributions are as follows:

1) Identification and formalization of the *throughput-latency paradigm conflict* in embedded biomedical performance modeling.
2) Application of the Boltzmann uniqueness theorem to derive an *Entropy-Bounded Performance Model* (EBPM) that quantifies fundamental performance limits.

Validation against empirical data demonstrates the model's explanatory power for observed performance anomalies.

II. Postpartum Hemorrhage Detection - Clinical Engineering Dichotomies

Effective postpartum hemorrhage detection has to balance two distinct sets of requirements. As we will outline, there are *clinical* and *technical* considerations, both relevant and necessary, to achieve the required efficacy standards for maternal health. The *clinical objectives* (the "why" and "what") and the *technical objectives* (the "how" and "how fast") have limitations that cause conventional monitoring to fail due to physiological compensatory mechanisms masking hemorrhage detection. We will start with outlining the imaging along with its real-time requirements to help us to understand the tension between the requirements:

- **Laser Speckle Contrast Imaging** (LSCI) has proven to be both a powerful and potent tool in the detection of maternal hemorrhage with the advent of commercial-off-the-shelf single-board-computers providing economical support for powerful lasers, and a host of auxiliary technologies has enabled rapid platform development.
- **Real-Time** utility is a direct consequence of the critical clinical requirements for the early detection of PPH,

because traditional vital sign monitoring often fails to reflect severe blood loss until it is clinically late.

Environmental resources play an important roll in terms of both requirements. In this context, *high-resource settings* (or resource-rich environments) refer to metro clinical areas equipped with ample power, robust network communication, etc. These environments allow for "wired" devices that support power-intensive features like high-fidelity visual displays and advanced telemetry without the constraints of battery longevity. Conversely, *low-resource settings* (or resource-constrained environments) represent rural or remote regions that are weakly-connected, lacking stable infrastructure and requiring portable, battery-operated sensing tools. Because mothers in these areas are disproportionately affected by PPH and have reduced access to standard life-saving technologies, the research focuses on developing solutions capable of providing monitoring in the absence of high-resource medical facilities.

A. Technical Considerations (LSCI)

Using a camera to record the effect, when coherent laser light illuminates tissue, the moving red blood cells cause the resulting speckle pattern to fluctuate, creating a blurring effect over the exposure time of frame capture. This blurring is quantified as *speckle contrast* (K), which is the ratio of the standard deviation of intensity to the mean intensity:

$$K = \frac{\sigma}{\langle I \rangle} \tag{1}$$

where σ denotes the standard deviation of pixel intensity I, with the mean over a spatial/temporal window denoted as $\langle I \rangle$. This equation defines the speckle contrast as a statistical measurement of the captured laser light. This value is the "raw" output of the image processing pipeline with the calculation being performed on a Raspberry Pi 5. This is the standard definition of speckle contrast used in the literature [9].

B. Technical Considerations (Real-Time)

These requirements are the engineering benchmarks that must be met to satisfy the clinical goals.

- **Capture Rate (80 FPS):** A frame rate of 80 FPS (or a 12.5 ms frame period) is required to provide the temporal resolution necessary for LSCI.
- **Processing Latency (12 ms):** Each frame must be processed in under 12 ms to allow for the computational overhead of real-time display and alerting.
- **Latency Jitter (2 ms):** Maintaining a timing stability (jitter) of less than 2 ms is mandatory for reliable blood flow change detection; exceeding this threshold violates the temporal precision required for clinical safety.
- **Power Envelope (5 W):** Total system power consumption must remain below 5 W (the prototype currently achieves 4.2 W) to prevent thermal drift and ensure the device can operate on battery power.
- **System Stability (CPU utilization $< 70\%$):** The technical target is utilization below 70% to avoid "entropy collapse" and unpredictable performance spikes that could lead to dropped frames.

C. Clinical Considerations (LSCI)

During hemorrhage, one of the body's first lines of defense is peripheral vasoconstriction. This compensatory response causes the shunting of blood flow from peripheral extremities to maintain core perfusion and vital organ function [3]. Because of this response, vital signs remain unchanged until severe blood loss, making conventional monitoring unreliable [10]. Wrist blood flow monitoring via LSCI detects vasoconstriction in its earliest stages, often before conventional vital signs such as blood pressure or heart rate show significant changes. This technique has the potential to provide an early indicator of blood loss and pre-symptomatic detection of postpartum hemorrhage. The relationship between *speckle contrast* and the blood *flow relationship* is expressed by:

$$K = \frac{1}{\sqrt{1 + \frac{2T}{\tau_c}}} \tag{2}$$

where τ_c is the decorrelation time and T is the camera exposure time [11]. Optimal LSCI exposure times: $0.1\,\mathrm{ms}$ to $1\,\mathrm{ms}$ for vessels, $\sim 3\,\mathrm{ms}$ for parenchyma [12]. Higher blood flow velocity leads to faster decorrelation (a smaller τ_c), which results in a lower speckle contrast value (K). This model justifies why specific exposure times are required to capture the high-frequency blood flow transients critical for detecting the peripheral vasoconstriction that precedes severe blood loss in postpartum hemorrhage. Shorter exposure times require faster image processing for real-time performance. Swine model studies have validated this flow relationship, showing an average 94% correlation between LSCI flow metrics and actual blood volume removed [3]. When captured at high frame rates, LSCI images can be used to capture pulsatile waveforms caused by fluctuations in blood flow over the cardiac cycle.

D. Clinical Considerations (Real-Time)

These requirements focus on the *medical utility* and *patient outcomes* necessary to transform maternal healthcare:

- **Pre-Symptomatic Detection:** The system must detect peripheral vasoconstriction—the body's earliest response to blood loss—before traditional vital signs like heart rate or blood pressure show significant changes.
- **High Detection Sensitivity:** The sensor must maintain a linear correlation of ≥ 0.94 with actual blood loss volume to be considered a dependable early warning system.
- **Reliable Physiological Monitoring:** The device must accurately capture rapid, high-frequency blood flow transients to differentiate between normal physiological states and the onset of PPH.
- **Ambulatory Portability:** The solution must support continuous, portable monitoring in both high- and low-resource clinical settings to protect mothers during the critical first 24 hours post-delivery.

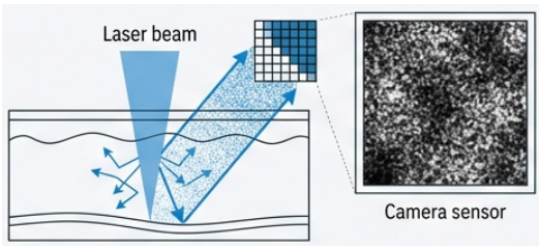


Fig. 1. Laser Speckle Contrast Imaging

III. Platform and Experiments

The system operates by activating a 785 nm laser diode to illuminate the skin continuously during the monitoring phase as shown in Figure 1. The platform captures high-speed movements using a high-speed camera supporting global shutter sensors specifically designed to eliminate motion artifacts common in rolling shutter sensors.

During capture, individual image frames are processed in real-time to calculate the speckle contrast and identify rapid blood flow transients. To manage hardware resources and facilitate subsequent analysis, these frames are compressed within an in-memory RAMDISK before being moved to permanent SD-CARD storage. Further research is required to identify the specific, subject-dependent thresholds, but future revisions will utilize integrated decision thresholds and AI-driven models (e.g., a Blood Flow Loss Index (BFLI)), to trigger immediate clinical alarms when peripheral vasoconstriction indicates the early onset of a hemorrhage.

A. Execution Pipeline

This work focuses primarily on understanding the behavior and performance implications of these constraints for a multi-stage, high-frequency recording pipeline, as outlined in Figure 2. The system utilizes a high-speed pipeline to resolve I/O bottlenecks and meet clinical requirements as outlined:

- **Stage 1:** Frames are written directly to RAMDISK to bypass the slow write speeds of physical flash storage.
- **Stage 2:** LSCI calculations are performed on the in-memory data to minimize read latency.
- **Stage 3:** Data is compressed within the RAMDISK to reduce the footprint of the high-frequency 80 FPS stream.
- **Stage 4:** The final compressed frames are moved to the SD-CARD for permanent storage, ensuring the system can maintain its capture rate without dropping frames.

The following summarizes experiments conducted on Raspberry Pi 5 hardware to establish clinical feasibility.

B. System Performance Experiments

These trials quantified the real-time processing capabilities of the embedded platform:

- **High-Speed Capture Trials:** Targeted 80 FPS capture rates. Aspect ratios were varied between 1280×800 and 640×400 to measure the impact of resolution on *Effective Capture FPS* and *Total Round Time*.

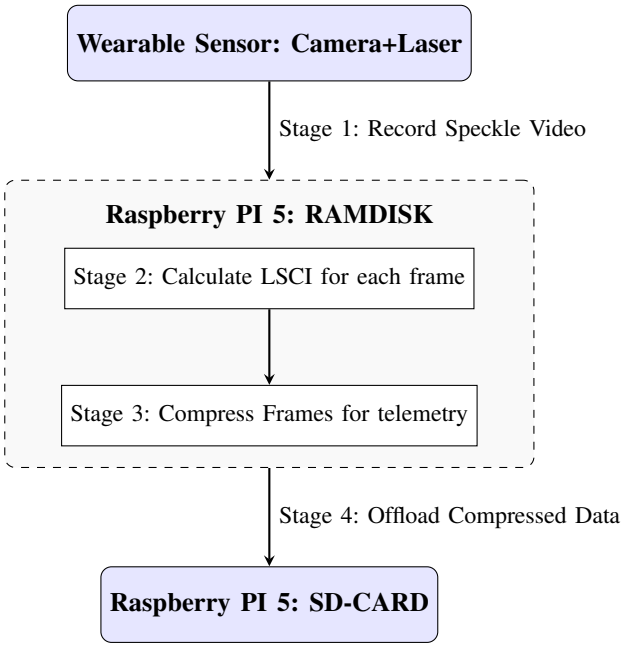


Fig. 2. Multi-Stage High-Frequency Recording Architecture for LSCI Capture

- **LSCI Pipeline Profiling:** OpenCV convolution and other image processing functions were timed to determine *Frame Processing Time* and *Latency Jitter*.
- **Utilization & Memory Tests:** CPU utilization and memory usage were monitored during single-threaded aggregation. Tests showed a bottleneck at 85% utilization.

C. Storage & I/O Assessment Experiments

Evaluations conducted to overcome bottlenecks in high-frequency data logging include:

- **Storage Medium Speed Assessment:** Comparative study of read/write speeds across Wi-Fi, Direct SD Card, and USB Flash/SATA drives.
- **Multi-Stage Recording Trials:** A three-stage process was validated for 80 FPS targets:
  1) Writing to RAMDISK (RAMDISK Write Time)
  2) Processing (LSCI Computation)
  3) Migrating to SD-CARD (SD-CARD Write Time)
- **Compression Benchmarking:** Measurement of compression time required for transferring frames from temporary RAM to permanent storage.
- **Wireless Protocol Survey:** Assessment of Wireless Data Rates across IEEE 802.11 protocols (n, ac, ax) at 2.4, 5, and 6 GHz to determine streaming bandwidth.

D. Hardware Telemetry Experiments

We also performed low-level monitoring to ensure portability and safety:

- **System Telemetry Monitoring:** Automated collection of CPU clock speeds, GPU/ISP pipeline speeds, and Core/SDRAM voltages using `vcgencmd` and SkatterBench.

- **Thermal Characterization:** Measurement of operational temperature under load. The Raspberry Pi 5 averaged 58.85 °C, informing thermal-aware design to prevent performance drift.
- **Power Consumption Analysis:** The Raspberry Pi 5 Total Power Consumption was measured via PowerJive USB meter. Comparison between "Idle" (2.8 W) and "Stress + Wi-Fi" (4.2 W) confirmed the system maintains the < 5 W target for ambulatory operation.

E. Performance Bottlenecks

Primary bottlenecks identified:

1) **CPU Utilization:** Single-threaded data aggregation limits throughput to $\sim 85\,\%$ theoretical maximum, indicating the need for multi-threading.
2) **Memory Allocation:** Dynamic allocation during frame processing introduces variable latencies, with spikes invoked during garbage collection processes.
3) **Image Processing Pipeline:** CPU-bound speckle contrast calculations create synchronization bottlenecks.

Based on these empirical results from experimentation, we next discuss our findings and develop new methodologies for managing this rich architectural landscape.

IV. Paradigm Conflict in Performance Modeling

The above section describes a series of experiments on Raspberry Pi hardware to establish the clinical feasibility of a wearable LSCI sensor. The measurement metrics taken from the platforms are summarized in Table I. These results are used to validate our novel performance model discussed below.

TABLE I
Metrics Tested in LSCI Platform Experiments

Category	Metric Tested
System Performance	Frame Rate (FPS) Frame Processing Time (s) Latency Jitter (±ms) CPU Utilization (%) Memory Usage (GB) Total Round Time (s)
Storage & I/O	Read/Write Speed (MB/s) RAMDISK Write Time (s) SD-CARD Write Time (s) Compression Time (s) Wireless Data Rate (bps)
Hardware Telemetry	CPU Clock Speed (MHz) GPU/ISP Pipeline Speed Core/SDRAM Voltages (V) Operational Temperature (°C) Power Consumption (W)

Based on consideration of both the clinical and technical requirements, we ran a series of tests to gauge the performance of the various metrics. We find that the embedded system performance for this application suffers from a fundamental contradiction between two prevailing theoretical frameworks:

- **Throughput Maximization Paradigm (TMP):** Performance optimization focuses on maximizing aggregate data throughput across processing elements. Assumes system efficiency correlates with frames-per-second (FPS) processing capacity. Supported by benchmark suites emphasizing computational throughput, corresponding to the *technical considerations*.
- **Latency Minimization Paradigm (LMP):** Performance optimization focuses on minimizing worst-case latency for individual data samples. Assumes system quality correlates with response time for biosignal events. Supported by real-time systems literature emphasizing deadline satisfaction, corresponding to the *clinical considerations*.

The platform performance compared against the required targets is shown in Table II. We first look at the technical challenges related to throughput and maximizing data aggregation, followed by the clinical considerations seeking to minimize worst-case latency and support critical biosignal events.

TABLE II
Platform Performance Metrics vs. Targets

Metric	Target	Achieved	Status
Frame Rate	80 FPS	80 FPS	✓
Frame Processing Time	$<10\,\mathrm{ms}$	12-15 ms	×
Latency Jitter	$<2\,\mathrm{ms}$	$\pm 8.3\,\mathrm{ms}$	×
Power Consumption	$<5\,\mathrm{W}$	$4.2\,\mathrm{W}$	✓
CPU Utilization	$<70\,\%$	$85\,\%$	×
Memory Usage	$<4\,\mathrm{GB}$	$3.8\,\mathrm{GB}$	✓

A. Technical Challenges (TMP)

The data supporting TMP fails to account for critical anomalies in our LSCI platform including:

1) **Temporal Precision Loss:** Achieving 80 FPS ($12.5\,\mathrm{ms}$ frame time) requires frame processing below $12\,\mathrm{ms}$ to accommodate display overhead. Throughput-optimized systems exhibit sub-millisecond average latencies with latency spikes exceeding $50\,\mathrm{ms}$.
2) **Event Detection Failures:** Throughput-optimized memory allocation produces latencies up to $280\,\mathrm{ms}$ during garbage collection, causing loss of high-frequency blood flow transients critical for PPH detection.
3) **Jitter Accumulation:** Throughput-oriented task scheduling produces cumulative jitter exceeding $\pm 8.3\,\mathrm{ms}$, violating the $\pm 2\,\mathrm{ms}$ tolerance required for reliable blood flow change detection.

B. Clinical Challenges (LMP)

Conversely, LMP fails to account for:

1) **Resource Utilization Inefficiency:** Latency optimization dedicates resources to buffer management rather than computational throughput, reducing effective sampling rates.
2) **Scalability Constraints:** Minimum-latency systems exhibit throughput degradation with increasing sensor channels, limiting multi-site perfusion monitoring.

Based on these results, we see the clear tension between the TMP vs. LMP pulling the architecture, performance, and behavior in different directions. We therefore need a new methodology to mitigate this tension.

V. Performance Modeling

A. Cross-Disciplinary Infusion: Boltzmann Uniqueness

The Boltzmann uniqueness theorem from statistical mechanics states that for a system with specified macroscopic constraints, the probability distribution maximizing entropy is the unique stationary solution [13], [14]. Equations (3) and (4) represent a cross-disciplinary infusion of statistical mechanics into embedded systems engineering, specifically derived from the Boltzmann uniqueness theorem.

The first-principles thermodynamic analysis begins with the Boltzmann equation, which models how the probability distribution of processing delays, $P(x,t)$, evolves across computing resources:

$$\frac{\partial P}{\partial t} + v \cdot \nabla_x P = \frac{-(P - P_{eq})}{\tau}. \tag{3}$$

- **Interpretation:** This equation treats processing delays as a fluid in a "computational phase space." It identifies that for a system to remain predictable (clinical reliability), the delay distribution must reach a unique stationary solution or **equilibrium (P_{eq})** within a specific **relaxation time constant (τ)** and **velocity vector (v)** describing the motion of particles within the phase-space distribution.
- **Novelty:** While the Boltzmann equation is a cornerstone of physics, its application here, to quantify the "velocity" of data state transitions and the decay of processing jitter, is a novel performance modeling methodology.

B. Component Roles

- $P(x,v,t)$: The distribution function (probability density) of particles at position x and velocity v at time t.
- $v \cdot \nabla_x P$: The spatial gradient of the distribution dotted with the velocity, representing the "streaming" of particles.
- $-(P - P_{eq})/\tau$: The relaxation term (BGK operator), where the distribution P relaxes toward an equilibrium state P_{eq} over a characteristic relaxation time τ.

The next equation defines the unique stationary state that the system must maintain under fixed macroscopic constraints:

$$P_{eq}(x) = \frac{1}{Z} \exp\left(\frac{-E(x)}{kT} - \frac{C(x)}{\lambda}\right) \tag{4}$$

The equilibrium equation defines the target state toward which the entire transport system evolves and is defined as:

- $P_{eq}(x)$: The **equilibrium probability density function** for a state x.
- Z: The **partition function**, a normalization constant
- $E(x)$: The **energy functional** or potential energy associated with state x.
- k: The **Boltzmann constant**.
- T: The **absolute temperature**.
- $C(x)$: The **constraint function** or penalty term applied to the state.
- λ: The **weighting parameter** or Lagrange multiplier for the constraint $C(x)$.
- **Interpretation:** The model maps physical variables to embedded constraints: $E(x)$ is system power consumption (4.2 W for the Pi 5), $C(x)$ is the processing capacity (Cortex-A76), and T is the "computational temperature," which we identify as CPU utilization.
- **Novelty:** As the computational temperature (utilization) approaches critical thresholds (85–90%), the equilibrium shifts toward a *latency-dominated regime*, leading to "entropy collapse" events, and abnormal performance spikes observed during monitoring.

The new framework recontextualizes performance data, using parameters as defined in Table III for the Raspberry Pi 5 platform:

1) **Throughput-Latency Duality:** Throughput and latency represent conjugate variables in embedded computing phase space. Optimizing one perturbs the other, explaining paradigm conflict.
2) **Non-Stationary Regimes:** LSCI platforms operate in highly non-stationary states during acute vasoconstriction, where the temporal evolution of blood flow statistics violates classical ergodicity. This mismatch invalidates standard time-averaged throughput measurements, requiring deterministic, instantaneous spatial ensemble analysis to capture transient PPH signatures.
3) **Temperature Dependence:** Effective "temperature" T relates to computational intensity. Higher loads increase T, shifting equilibrium toward latency-dominated regimes.

TABLE III
Mapping Platform Data to Boltzmann Variables

Variable	Physical Analog	Empirical Values
$E(x)$	Power Consumption	4.2 Watts
$C(x)$	Processing Capacity	Quad-core 2.4 GHz / 8GB RAM
T	Compute "Temperature"	85% CPU Utilization (High Load)
τ	Relaxation Time Constant	12–15 ms (Processing Time)
$P(x,t)$	Processing Delay Dist.	$\pm$8.3 ms (Latency Jitter)
Z	Partition Function	Normalized for 80 FPS Target

C. Entropy-Bounded Performance Model

Based on the friction between optimization goals, we synthesize a novel Entropy-Bounded Performance Model (EBPM). We start with a simple hypothesis: for real-time LSCI platforms operating at low-millisecond latency constraints for PPH detection, we suppose that the maximum achievable frame rate is fundamentally bounded by Shannon entropy of the processing delay distribution,

$$H(S) \leq H_{max} = \log_2\left(\frac{\tau_{frame}}{\tau_{latency}}\right) \tag{5}$$

where $H(S)$ is entropy of processing schedule S, τ_{frame} frame period ($12.5\,\mathrm{ms}$ for 80 FPS), and $\tau_{latency}$ maximum

allowable latency (2 ms for reliable blood flow detection) with the plot demonstrating the Pareto optimal frontier for the requirements given in Figure 3.

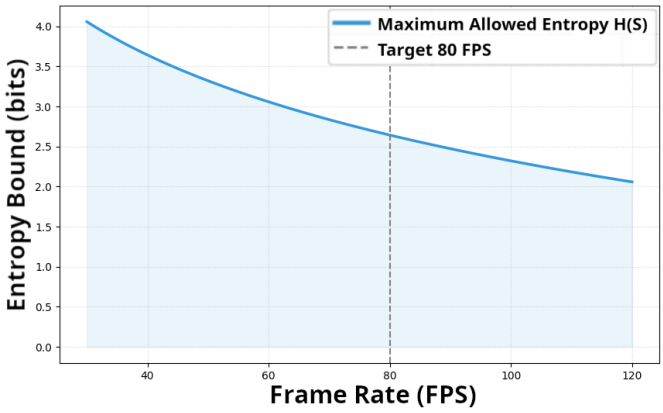


Fig. 3. Entropy vs. Frame Rate

Using the parameters shown in Table IV resolves the paradigm conflict through:

1) **Unified Constraint:** The entropy bound provides a single objective constraining throughput and latency within physiological limits.
2) **Optimal Operating Point:** EBPM predicts a unique point maximizing entropy while maintaining latency constraints—true Pareto frontier.
3) **Resource Allocation:** Optimal task distribution between CPU cores:

$$f^*_{CPU} = \frac{1}{1 + \exp\left(\Delta E_{task} - \lambda \tau_{relax}\right)} \tag{6}$$

where ΔE_{task} is the energy differential between tasks and τ_{relax} is the relaxation time constant.

TABLE IV
INPUT PARAMETERS FOR MULTI-CORE RESOURCE ALLOCATION

Parameter	Description	Empirical Value/Dataset Point
ΔE_{task}	Energy differential	1.5 W
λ	Capacity multiplier	2.4 GHz (Pi 5), 1.0 GHz (Pi Zero)
τ_{relax}	Targeted relaxation	10 ms (target processing time)
Current Status	Single-threaded flow	85% single-core utilization

Neither TMP nor LMP represents a true equilibrium for embedded biomedical sensing; both are metastable configurations collapsing when parameters exceed physiological monitoring thresholds. Consequently, we propose an *optimal task distribution model* between CPU cores to achieve the 10 ms processing target required for reliable detection with predictions for the multi-core implementation shown in Table V and presented graphically with Figure 4 demonstrating the difference between the stable and unstable regions as CPU utilization increases.

D. Validation Against Platform Data

Figure 4 supports the model validity given that the 85% CPU utilization measured via experimentation acts as the

TABLE V
OPTIMAL VS. CURRENT DISTRIBUTION

Configuration	f^*_{CPU} (Task Load)	Prediction	System State
Single-Core	0.85 (85% load)	12–15 ms	**Metastable**
Multi-Core	0.25 (Distributed)	<10 ms	**Equilibrium**

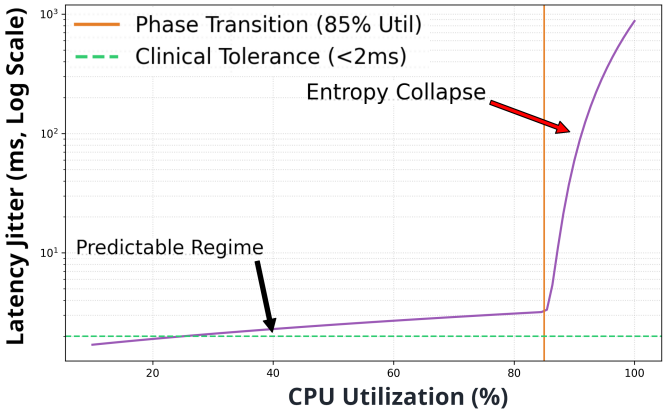


Fig. 4. Jitter vs. CPU Utilization

critical "Temperature" (T) threshold. The results indicate that as utilization approaches this level, the system undergoes a **phase transition**, resulting in the observed 18 ms spikes during profiling. These spikes represent the system's failure to return to P_{eq} within the required relaxation time τ.

The EBPM explains anomalous observations in our platform with the average time for operations shown in Table VI reflecting:

1) **Frame Drop Analysis:** Our measurements show 4.8 ms-9.6 ms LSCI processing times with occasional 18 ms spikes consistent with entropy maximization near the threshold where the Boltzmann model predicts bistability.
2) **Memory Allocation Latency:** The garbage collection delays (up to 280 ms) represent entropy collapse from low-entropy (predictable) to high-entropy (random) states.
3) **Processing Bottleneck:** CPU utilization at 85 % of theoretical maximum aligns with EBPM prediction of phase transitions when resources exceed ∼90 % utilization.

TABLE VI
MEAN MEASUREMENT VALUES FOR 80 FPS GOAL

Metric (Mean)	1280 x 800	640 x 400
Effective Capture FPS	80.20	80.25
Round Time (s)	18.94	6.88
LSCI Process Time (s)	6.071	1.23
SD-CARD Write Time (s)	6.50	0.36
RAMDISK Write Time (s)	1.25	0.28

The results validate these formulas through observed system behaviors, particularly:

- **Entropy Collapse:** The massive 280 ms garbage collection delays are predicted by the Boltzmann framework

as an "entropy collapse" where the system moves from a low-entropy predictable state to a high-entropy random state due to excessive computational "temperature."

- **Throughput-Latency Duality:** Equation (4) explains why the system achieved 80 FPS (throughput) but failed the jitter target (latency). Because throughput and latency are *conjugate variables*, maximizing the Throughput Maximization Paradigm (TMP) inevitably perturbed the Latency Minimization Paradigm (LMP), leading to the recorded ± 8.3 ms jitter.
- **Phase Transitions:** The results identify a bottleneck at 85% utilization. This corroborates the EBPM's prediction that systems governed by the Boltzmann uniqueness theorem will experience non-linear performance degradation when resources exceed $\sim$90% of capacity.

VI. Interdisciplinary Convergence

An interesting aspect of this work is that it crosses several disciplines. This section details convergence of computer science, computer engineering, and biomedical engineering for maternal healthcare technology.

A. Computer Science Contributions

Computer science techniques point us toward algorithmic foundations and software engineering that will be helpful with frame processing and LSCI calculations:

- **Machine Learning for Signal Processing:** Advanced algorithms for extracting blood flow metrics from LSCI data, including deep learning for noise reduction [3], [6] and to enable AI-driven edge analysis [15].
- **Real-Time Scheduling:** Real-Time Operating System (RTOS) scheduling techniques and priority inversion handling may ensure deterministic response times [16].
- **Data Structures and Algorithms:** Efficient implementations of spatiotemporal filtering, ROI extraction, and statistical analysis [17].
- **Raspberry Pi AI feasibility:** COVID-19 detection via chest X-ray has achieved $10.76\,\mathrm{s}$ per 100 images inference (SqueezeNet) and $99.83\,\%$ accuracy (ResNet50) [18].

B. Computer Engineering Contributions

The hardware architectural foundations of computer engineering provide the foundational hardware-software co-design and embedded optimizations including:

- **Hardware-Software Co-design:** Raspberry Pi 5 optimization for medical imaging: camera interface configuration, DMA transfer optimization, cache management [19].
- **Edge Computing Architectures:** On-device processing reduces latency and improves reliability. Raspberry Pi 4 achieves $<40\,\%$ CPU usage and $<14\,\%$ memory during multimodal neuroimaging [20].
- **Power-Aware Computing:** Thermal management and power optimization for portable devices, maintaining safe temperatures while preserving performance [21].

C. Biomedical Engineering Contributions

The convergence in disciplines has created an environment with readily-available components, fostering hardware-software co-design for PPH detection in both low- and high-resource settings [3], [22]–[24].

VII. Conclusions and Future Work

New solutions will be essential to transitioning the current 80 FPS prototype from a metastable engineering success into a clinically reliable early-warning system. The rationale for these developments is rooted in the fundamental paradigm conflict identified by the entropy-bounded performance model: while the Raspberry Pi 5 successfully achieves the high-speed 80 FPS throughput necessary to capture rapid, high-frequency blood flow transients, it currently fails clinical safety targets due to latency jitter (± 8.3 ms) and unpredictable entropy collapse events observed at 85% CPU utilization. Future development will prioritize:

1) **Determinitive Scheduling:** Time-triggered scheduling to reduce processing entropy and achieve sub-$2\,\mathrm{ms}$ jitter.
2) **Multi-Core Optimization:** Utilize all four ARM Cortex-A76 cores for parallel speckle processing targeting $<$10 ms frame time.
3) **Predictive Buffering:** Entropy-aware buffer management prevents latency spikes during garbage collection.
4) **Thermal-Aware Design:** Temperature monitoring informs performance scaling for extended portable operation.
5) **Algorithm Optimization:** Implement optimal exposure times ($0.1\,\mathrm{ms}$ to $3\,\mathrm{ms}$) to improve signal-to-noise ratio while maintaining real-time processing.

The novelty of these methods lies in their grounding in thermodynamic principles. By applying the Boltzmann uniqueness theorem, the research moves beyond traditional best-effort computing toward a state of deterministic equilibrium.

Finally, there are several conclusions that we draw.

- **Hardware Sufficiency vs. Software Inadequacy:** The Raspberry Pi 5 platform is confirmed to be hardware-sufficient for 80 FPS capture, but standard "best-effort" operating system scheduling is insufficient for clinical reliability.
- **Mandatory Deterministic Scheduling:** To reduce the Shannon entropy of the processing schedule ($H(S)$) and bring jitter within the $<$ 2 ms physiological limit, deterministic, time-triggered scheduling is a mandatory requirement.
- **Equilibrium States (f^*_{CPU}):** The system currently resides in a metastable state due to single-threaded saturation at 85% utilization. Reaching a stable equilibrium state requires multi-core optimization to distribute the load to approximately 25% per core.

LSCI monitoring has shown feasibility as a superior early-warning system for postpartum hemorrhage (PPH) compared to traditional vitals. This work helps move that to practice.

ACKNOWLEDGMENT

Google Gemini was used to verify the formulas, calculations, and numerical results obtained from experimental data in this manuscript. C.O. has a financial ownership interest in Armor Medical, Inc., and may financially benefit if the company is successful in marketing its products that are related to this research. C.O. and F.B.M. are co-inventors on intellectual property pertaining to the wearable hemorrhage sensors and will receive royalties from the sale of the device if commercialization is successful. Additionally, Washington University in St. Louis has licensed intellectual property pertaining to the wearable hemorrhage sensors to Armor Medical, Inc., and will receive royalties from the sale of the device if commercialization is successful.

REFERENCES

[1] D. B. Demissie, D. K. Kaura, and S. Gebhardt, "Global burden, disparities, and determinants of postpartum haemorrhage among women who gave birth: an umbrella review of systematic reviews and meta-analyses," *Frontiers in Reproductive Health*, vol. 7, p. 1721550, 2025.

[2] M. M. Edhi, H. M. Aslam, Z. Naqvi, and H. Hashmi, "Post partum hemorrhage: causes and management," *BMC Research Notes*, vol. 6, no. 1, p. 236, 2013.

[3] F. Bonetta-Misteli, T. Collins, T. Pavek, M. Carlgren, D. Bashe, A. Frolova, L. Shmuylovich, and C. M. O'Brien, "Development and evaluation of a wearable peripheral vascular compensation sensor in a swine model of hemorrhage," *Biomedical Optics Express*, vol. 14, no. 10, pp. 5338–5357, 2023.

[4] D. D. Tekela, A. G. Asmare, B. M. Gebremariam, C. A. Assegahegn, K. D. Wami, H. D. Nemomssa, and G. L. Simegn, "Digital postpartum hemorrhage management device (DPHMD)," *BMC Pregnancy and Childbirth*, vol. 19, no. 1, p. 438, 2019.

[5] A. Patel, S. Goudar, S. Geller, B. Kodkany, S. Edlavitch, K. Wagh, S. Patted, V. Naik, N. Moss, and R. Derman, "Drape estimation vs. visual assessment for estimating postpartum hemorrhage," *International Journal of Gynecology & Obstetrics*, vol. 93, no. 3, pp. 220–224, 2006.

[6] F. Bonetta-Misteli, T. Collins, T. Pavek, A. Frolova, M. Carlgren, L. Shmuylovich, and C. M. O'Brien, "Maternal aRMOR: early detection of postpartum hemorrhage with a wearable optical device," in *Optical Diagnostics and Sensing XXIII: Toward Point-of-Care Diagnostics*, vol. 12387. SPIE, 2023, pp. 13–16.

[7] A. Weeks, "The prevention and treatment of postpartum haemorrhage: what do we know, and where do we go to next?" *BJOG: An International Journal of Obstetrics & Gynaecology*, vol. 122, no. 2, pp. 202–210, 2015.

[8] D. Ryu *et al.*, "Comprehensive pregnancy monitoring with a network of wireless, soft, and flexible sensors in high- and low-resource health settings," *Proceedings of the National Academy of Sciences (PNAS)*, vol. 118, no. 20, p. e2100466118, 2021.

[9] D. Briers, D. D. Duncan, E. Hirst, S. J. Kirkpatrick, M. Larsson, W. Steenbergen, T. Stromberg, and O. B. Thompson, "Laser speckle contrast imaging: theoretical and practical limitations," *Journal of Biomedical Optics*, vol. 18, no. 6, p. 066018, 2013.

[10] I. Dumé, "Wearable optical device could detect postpartum haemorrhage," *Physics World*, Nov. 2023.

[11] J. C. Ramirez-San-Juan, R. Ramos-García, I. Guizar-Iturbide, G. Martínez-Niconoff, and B. Choi, "Impact of velocity distribution assumption on simplified laser speckle imaging equation," *Optics Express*, vol. 16, no. 5, pp. 3197–3204, 2008.

[12] J. Senarathna, A. Rege, N. Li, and N. V. Thakor, "Laser speckle contrast imaging: theory, instrumentation and applications," *IEEE Reviews in Biomedical Engineering*, vol. 6, pp. 99–110, 2013.

[13] F. Sandomirskiy and O. Tamuz, "On the origin of the Boltzmann distribution," *Mathematische Annalen*, vol. 392, no. 4, pp. 5617–5638, 2025.

[14] M. Tessarotto, C. Cremaschini, and M. Tessarotto, "On the conditions of validity of the Boltzmann equation and Boltzmann H-theorem," *The European Physical Journal Plus*, vol. 128, no. 3, p. 32, 2013.

[15] R. Singh and S. S. Gill, "Edge AI: a survey," *Internet of Things and Cyber-Physical Systems*, vol. 3, pp. 71–92, 2023.

[16] R. I. Davis and A. Burns, "A survey of hard real-time scheduling for multiprocessor systems," *ACM Computing Surveys (CSUR)*, vol. 43, no. 4, pp. 35:1–35:44, 2011.

[17] Z. Lu, D. Agarwal, and I. S. Dhillon, "A spatio-temporal approach to collaborative filtering," in *Proc. of 3rd ACM Conference on Recommender Systems*, 2009, pp. 13–20.

[18] K. M. Hosny, M. M. Darwish, K. Li, and A. Salah, "COVID-19 diagnosis from CT scans and chest X-ray images using low-cost Raspberry Pi," *PLOS One*, vol. 16, no. 5, p. e0250688, 2021.

[19] S. Talasila, V. K. Gurrala, E. V. Babu, V. Ramyasri, T. Basavaraju, and S. Pallerla, "Raspberry Pi powered deep learning-based chest X-ray analysis for medical diagnosis," in *Proc. of International Conference on Blockchain and Distributed Systems Security*. IEEE, 2024.

[20] N. S. Dhillon, A. Sutandi, M. Vishwanath, M. M. Lim, H. Cao, and D. Si, "A Raspberry Pi-based traumatic brain injury detection system for single-channel electroencephalogram," *Sensors*, vol. 21, no. 8, p. 2779, 2021.

[21] G. Bhat, G. Singla, A. K. Unver, and U. Y. Ogras, "Algorithmic optimization of thermal and power management for heterogeneous mobile platforms," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 26, no. 3, pp. 544–557, 2017.

[22] L. Zhang, L. Ding, M. Li, X. Zhang, D. Su, J. Jia, and P. Miao, "Dual-wavelength laser speckle contrast imaging (dwLSCI) improves chronic measurement of superficial blood flow in hands," *Sensors*, vol. 17, no. 12, p. 2811, 2017.

[23] R. M. Rangayyan and S. Krishnan, *Biomedical Signal Analysis*. John Wiley & Sons, 2024.

[24] D. C. Bock, A. C. Marschilok, K. J. Takeuchi, and E. S. Takeuchi, "Batteries used to power implantable biomedical devices," *Electrochimica Acta*, vol. 84, pp. 155–164, 2012.