

# Final Report: Large-Scale N-Body Accelerator

Lucy He (lh3365), Xiyuan Peng (xp2236), Jingzeng Xie (jx2668),  
Pengpeng Wang (pw2660), Charlotte Chen (hc3558)

Spring 2026

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>System Block Diagram</b>	<b>1</b>
<b>3</b>	<b>The N-Body Problem</b>	<b>2</b>
<b>4</b>	<b>Hardware Implementation</b>	<b>3</b>
4.1	Key Computation Building Blocks . . . . .	3
4.1.1	Floating Point Adder . . . . .	3
4.1.2	Floating Point Multiplier . . . . .	4
4.1.3	Fast Inverse Square Root . . . . .	4
4.1.4	Pipelined Two-Body Core . . . . .	6
4.2	N-body Memory . . . . .	6
4.3	Four-Core Wrapper and Data Flow . . . . .	7
4.3.1	Tiled Input Organization . . . . .	7
4.3.2	Broadcasted Source Body . . . . .	7
4.3.3	Accumulation and Feedback . . . . .	8
4.3.4	Output Selection and Writeback . . . . .	8
4.4	Leapfrog Integrator . . . . .	8
4.4.1	Leapfrog Integration Algorithm . . . . .	8
4.4.2	Integrator Implementation . . . . .	9
4.5	N-body Controller . . . . .	9
<b>5</b>	<b>Software Implementation</b>	<b>11</b>
5.1	Software Architecture . . . . .	12
5.2	Real-Time Multi-Threaded Software Design . . . . .	12
5.3	Major Software Components . . . . .	13
5.4	Shared Simulation State and History Buffer . . . . .	13
5.5	Initialization and Reset Flow . . . . .	14
5.6	Accelerator Thread . . . . .	14
5.7	Display Thread and Frame Rendering . . . . .	16
5.7.1	Precomputed Circular Body Masks . . . . .	16
5.8	Keyboard Input and User Interface . . . . .	18
5.9	Kernel Driver Layer . . . . .	18
5.10	End-to-End Software Flow . . . . .	19
<b>6</b>	<b>Hardware/Software Interface</b>	<b>19</b>
6.1	Accelerator Interface . . . . .	19
6.2	Display Interface . . . . .	19

<b>7</b>	<b>Testbench and Verification</b>	<b>21</b>
7.1	Arithmetic Unit Testbenches . . . . .	21
7.1.1	Floating Pint Adder . . . . .	21
7.1.2	Floating Pint Multiplier . . . . .	21
7.1.3	Fast Inverse Square Root . . . . .	21
7.1.4	N-body Memory . . . . .	21
7.2	Golden Model . . . . .	22
7.3	Frame Data Generation . . . . .	23
7.4	Outputs and Comparison . . . . .	23
7.4.1	Acceleration Accumulation Outputs . . . . .	23
7.4.2	Controller-Integrated Outputs . . . . .	23
7.4.3	Software-visible Outputs . . . . .	24
<b>8</b>	<b>Performance and Resource Utilization</b>	<b>25</b>
8.1	Performance Metrics . . . . .	25
8.2	Post-Fit FPGA Resource Utilization . . . . .	26
8.2.1	Custom 27-bit Floating-Point Format . . . . .	27
8.2.2	DSP Usage by Arithmetic Submodule . . . . .	27
8.2.3	Memory Usage . . . . .	28
<b>9</b>	<b>Contributions and Lessons Learned</b>	<b>29</b>
<b>A</b>	<b>Code</b>	<b>31</b>
A.1	Software Source Code . . . . .	31
A.1.1	main.c . . . . .	31
A.1.2	nbody.c . . . . .	32
A.1.3	nbody.h . . . . .	37
A.1.4	display.c . . . . .	38
A.1.5	display.h . . . . .	42
A.1.6	keyboard.c . . . . .	42
A.1.7	keyboard.h . . . . .	45
A.1.8	usbkeyboard.c . . . . .	45
A.1.9	usbkeyboard.h . . . . .	46
A.1.10	accelerator_driver.c . . . . .	47
A.1.11	nbody_ioctl.h . . . . .	52
A.1.12	display_driver.c . . . . .	53
A.1.13	display_ioctl.h . . . . .	55
A.1.14	Makefile . . . . .	55
A.2	Software Test Code . . . . .	56
A.2.1	avmm_smoke_test.c . . . . .	56

A.2.2	avmm_frame_xy_dump.c . . . . .	60
A.2.3	Test Makefile . . . . .	66
A.3	Hardware Source Code . . . . .	67
A.3.1	nbody_accel_avmm.sv . . . . .	67
A.3.2	vga_bitmap_avmm.sv . . . . .	70
A.3.3	nbody_control.sv . . . . .	73
A.3.4	four_core_wrapper.sv . . . . .	81
A.3.5	fourcore_bcj_datapath.sv . . . . .	84
A.3.6	two_body_core.sv . . . . .	86
A.3.7	nbody_integrator.sv . . . . .	91
A.3.8	nbody_mem.sv . . . . .	93
A.3.9	framebuffer_ram.sv . . . . .	95
A.3.10	FpAdd.sv . . . . .	95
A.3.11	FpMul.sv . . . . .	97
A.3.12	FpInvSqrt.sv . . . . .	98
A.3.13	FpNegate.sv . . . . .	100
A.4	Hardware Testbenches . . . . .	100
A.4.1	tb_nbody_control.sv . . . . .	100
A.4.2	tb_four_core_wrapper.sv . . . . .	105
A.4.3	tb_core_accel.sv . . . . .	110
A.4.4	tb_nbody_integrator.sv . . . . .	114
A.4.5	tb_nbody_mem.sv . . . . .	117
A.4.6	tb_fpadd.sv . . . . .	121
A.4.7	tb_fpmul.sv . . . . .	122
A.4.8	tb_fpinvsqrt.sv . . . . .	123
A.5	Golden Model and Input Generation Code . . . . .	125
A.5.1	golden.py . . . . .	125
A.5.2	golden_control.py . . . . .	139
A.5.3	golden_avmm_xy.py . . . . .	144
A.5.4	FP64.py . . . . .	146
A.5.5	input_gen_s1e8m18.py . . . . .	149

# 1 Introduction

The goal of this project is to design and implement an FPGA-based hardware accelerator for gravitational N-body simulation, a classic  $O(N^2)$  problem in which each body interacts with every other body through pairwise forces (1). Our system targets the force-computation and time-update loop of a 2D gravity simulation. The hardware consists of four parallel two-body compute cores with pipelined datapath. Each core computes pairwise acceleration contributions using a fast inverse square root unit, followed by accumulation. A leapfrog integration step then updates particle velocities and positions over time. The host processor handles initialization, user interaction, and display coordination via an Avalon memory-mapped interface. The resulting simulation is visualized on a VGA display as an interactive 2D gravity simulator.




Figure 1: System Demo

# 2 System Block Diagram

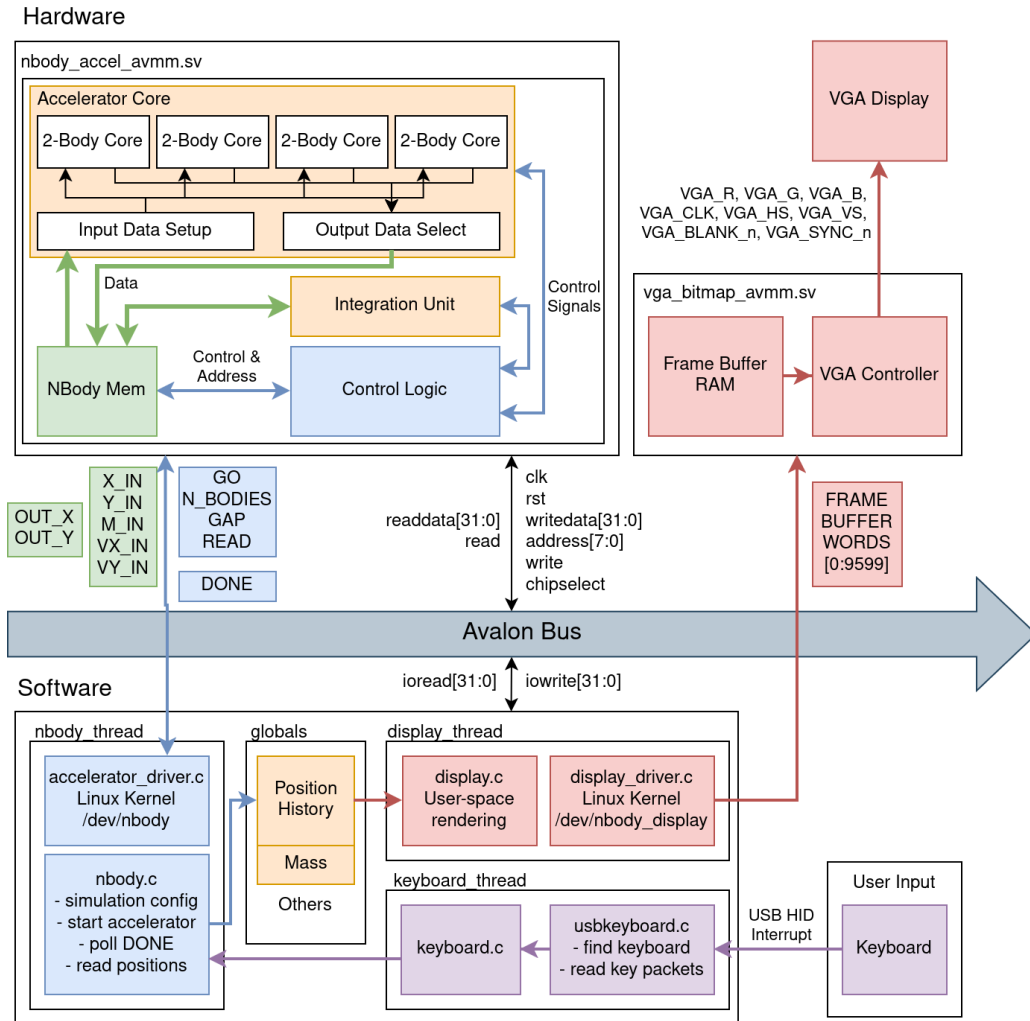


Figure 2: System Block Diagram

Figure 2 shows the full hardware/software organization of the N-body accelerator system. The design is split into two Avalon-MM peripherals: the N-body accelerator module, `nbody_accel_avmm.sv`, and the VGA display module, `vga_bitmap_avmm.sv`. Both peripherals are connected to the HPS software through the Avalon bus and are accessed from software using memory-mapped I/O through Linux kernel drivers.

The accelerator peripheral contains local N-body memory, four parallel 2-body computation cores, an integration unit, and control logic. Software writes body data and configuration registers through the Avalon bus, starts computation using the `GO` signal, and reads back updated positions when the hardware raises `DONE`.

The VGA peripheral contains a frame buffer RAM and VGA controller. Software converts body positions into pixels and writes the frame buffer through the Avalon bus. The VGA controller continuously reads this buffer and generates the VGA timing and RGB output signals for the monitor.

On the software side, the application uses separate threads for acceleration, display, and keyboard input. The accelerator thread controls the hardware simulation step, the display thread renders the current frame, and the keyboard thread handles user input. Shared global data stores the current body states and position history.

Overall, software manages simulation control and user interaction, while the FPGA hardware performs the expensive N-body computation and real-time VGA display. The Avalon bus connects the HPS software to both hardware peripherals.

### 3 The N-Body Problem

The N-body problem simulates particles whose motion is determined by pairwise interactions with all other bodies. In the direct formulation, each of the  $N$  bodies interacts with the other  $N - 1$  bodies, resulting in a computational complexity of  $O(N^2)$ .

The evolution of the system begins with the standard Newtonian gravity equation. For two bodies with masses  $m_i$  and  $m_j$ , the magnitude of the gravitational force is given by:

$$|F_{ij}| = G \frac{m_i m_j}{r_{ij}^2}, \quad (1)$$

where  $G$  is the gravitational constant and  $r_{ij}$  denotes the distance between the two bodies.

By applying Newton's second law, the acceleration induced by body  $j$  on body  $i$  can be written as:

$$a_{ij} = \frac{F_{ij}}{m_i} = G \frac{m_j}{r_{ij}^2}. \quad (2)$$

After summing the contributions from all other bodies, the simulation obtains the total acceleration vector of body  $i$ . To express both magnitude and direction in vector form, the relative displacement vector is introduced as  $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ . Accordingly, the acceleration contribution from body  $j$  to body  $i$  is written as:

$$\mathbf{a}_{ij} = G \frac{m_j \mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|^3}. \quad (3)$$

In practical simulation, a softening factor  $\varepsilon^2$  is added to avoid numerical instability when two bodies become too close. Therefore, the softened pairwise interaction is expressed as:

$$\mathbf{a}_{ij} \approx G \frac{m_j \mathbf{r}_{ij}}{(\|\mathbf{r}_{ij}\|^2 + \varepsilon^2)^{3/2}}. \quad (4)$$

For the two-dimensional implementation adopted in this work, the acceleration vector can be further decomposed. Let  $\Delta x_{ij} = x_i - x_j$  and  $\Delta y_{ij} = y_i - y_j$ , then the squared distance between bodies  $i$  and  $j$  is:

$$r_{ij}^2 = (\Delta x_{ij})^2 + (\Delta y_{ij})^2. \quad (5)$$

Accordingly, the softened pairwise acceleration can be expressed component-wise as:

$$a_{ij,x} \approx Gm_j \frac{\Delta x_{ij}}{(r_{ij}^2 + \varepsilon^2)^{3/2}}, \quad a_{ij,y} \approx Gm_j \frac{\Delta y_{ij}}{(r_{ij}^2 + \varepsilon^2)^{3/2}}. \quad (6)$$

The total acceleration of body  $i$  is then obtained by summing these contributions over all interacting bodies in the  $x$  and  $y$  directions separately. This decomposition is used directly in hardware, where the accelerator computes the two directional components before passing them to the integration stage.

The goal of this project is to accelerate the  $O(N^2)$  pairwise interaction computation, which is the main bottleneck in direct N-body simulation. To improve throughput, the hardware uses four parallel, fully pipelined 2-body computation cores. The design also adopts a custom 27-bit floating-point format so that each mantissa multiplication maps efficiently to one  $18 \times 18$  DSP block. In addition, a fast inverse square root (FISR) unit is used to efficiently approximate the inverse-distance term, and a leapfrog integration unit is implemented in hardware to update body velocity and position after acceleration computation.

## 4 Hardware Implementation

### 4.1 Key Computation Building Blocks

The accelerator is composed of several core computation modules that together implement the N-body pipeline. These hardware building blocks are designed around the custom 27-bit floating-point datapath and are optimized for efficient FPGA execution.

The main computation modules include:

- **Floating Point Adder** Performs floating-point accumulation and arithmetic reduction operations required during force summation.
- **Floating Point Multiplier** Computes the multiplication operations used in distance calculation, force scaling, and acceleration updates.
- **Fast Inverse Square Root** Implements an optimized approximation of  $1/\sqrt{x}$ , which is the computational bottleneck of the gravitational interaction equation.
- **Two Body Core** Computes the pairwise gravitational interaction between two particles and generates the corresponding acceleration contribution.

Together, these modules form the computational foundation of the accelerator and enable large-scale parallel N-body simulation on the FPGA platform.

#### 4.1.1 Floating Point Adder

The floating-point adder is implemented as a two-stage pipelined unit. In the first stage, the exponents of the two inputs are compared, and the mantissa with the smaller exponent is shifted to align both operands to the same exponent. In the second stage, the aligned mantissas are either added or

subtracted depending on the input signs, followed by normalization to produce the final floating-point result.

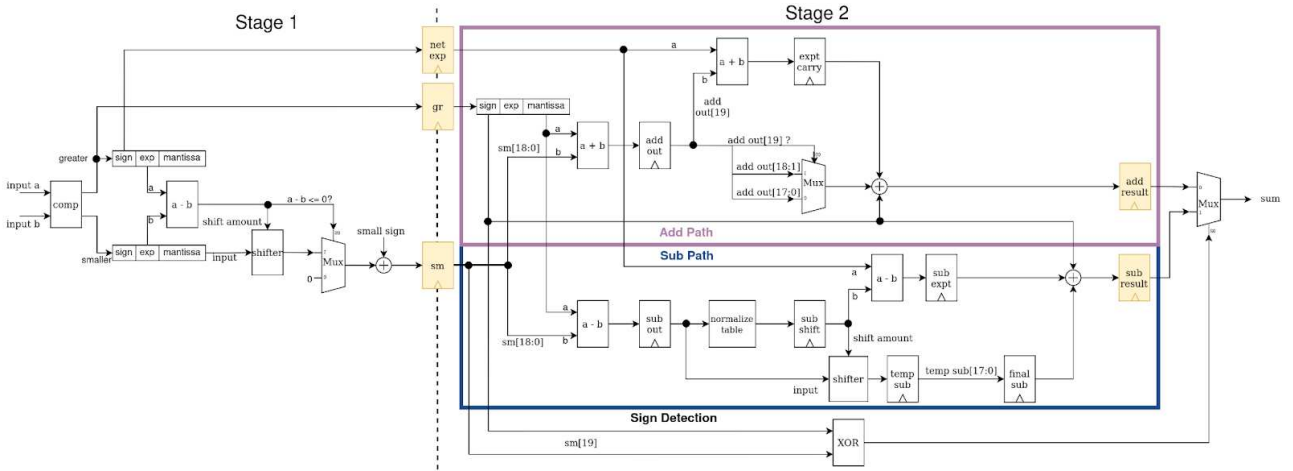


Figure 3: Fpadder Mudule

#### 4.1.2 Floating Point Multiplier

The floating-point multiplier is implemented as a fully combinational unit. The output sign is computed by XORing the two input signs, the output exponent is computed by adding the input exponents with bias correction, and the mantissas are multiplied to form the product. The result is then normalized and packed back into the custom floating-point format.

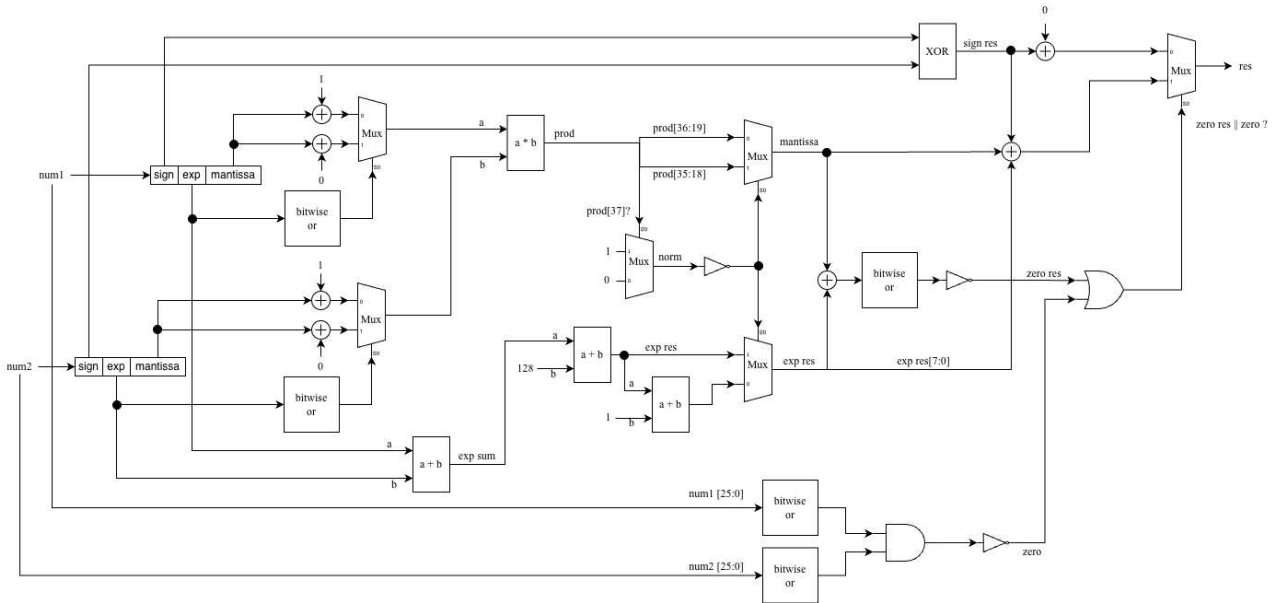


Figure 4: Fpmultiplier Mudule

#### 4.1.3 Fast Inverse Square Root

The core computational bottleneck in the N-body force calculation is evaluating the term:

$$g(r^2) = \frac{1}{(r^2 + \varepsilon^2)^{3/2}}, \quad (7)$$

which requires both a square root and a division. On general-purpose hardware, dedicated square root and division units are expensive in terms of area and latency, making them poorly suited to a

deeply pipelined FPGA design. Instead, we adopt the well-known Fast Inverse Square Root (FISR) algorithm, originally popularized by the Quake III Arena source code, which approximates  $1/\sqrt{x}$  using only integer bit manipulation, addition, and multiplication (2; 3).

**Step 1: Integer Interpretation of a Float.** The key insight is that the bit pattern of an IEEE 754 single-precision float encodes an approximation of its own base-2 logarithm. A 32-bit floating-point number is stored as three fields  $[S | E | M]$ , where  $S$  is the sign bit,  $E$  is the 8-bit biased exponent, and  $M$  is the 23-bit mantissa:

$$x = (-1)^S \cdot 2^{E-127} \cdot \left(1 + \frac{M}{2^{23}}\right).$$

Its integer reinterpretation is  $I = S \cdot 2^{31} + E \cdot 2^{23} + M$ , giving the approximation:

$$\log_2(x) \approx \frac{I}{2^{23}} - 127.$$

**Step 2: Magic Number.** To compute  $y \approx 1/\sqrt{x}$ , we require  $\log_2(y) \approx -\frac{1}{2} \log_2(x)$ . Substituting the integer approximation and solving for  $I_y$  gives:

$$I_y \approx c_0 - \frac{I_x}{2},$$

where the magic constant  $c_0 = 0x5f3759df$  accounts for the floating-point bias and a small correction term. The original Quake III implementation is:

```
float Q_rsqrt(float x) {
    float x2 = 0.5f * x;
    int i = *(int*)&x;
    i = 0x5f3759df - (i >> 1);
    float y = *(float*)&i;
    y = y * (1.5f - x2 * y * y);
    return y;
}
```

**Step 3: Newton's Method.** The initial estimate requires refinement using Newton–Raphson iteration (4). Setting  $f(y) = 1/y^2 - x = 0$ , whose root is  $y = 1/\sqrt{x}$ , with  $f'(y) = -2/y^3$ , one Newton–Raphson iteration gives:

$$y_{n+1} = y_n (1.5 - 0.5 \cdot x \cdot y_n^2). \quad (8)$$

This correction requires two multiplications and one subtraction with no division or square root hardware involved.

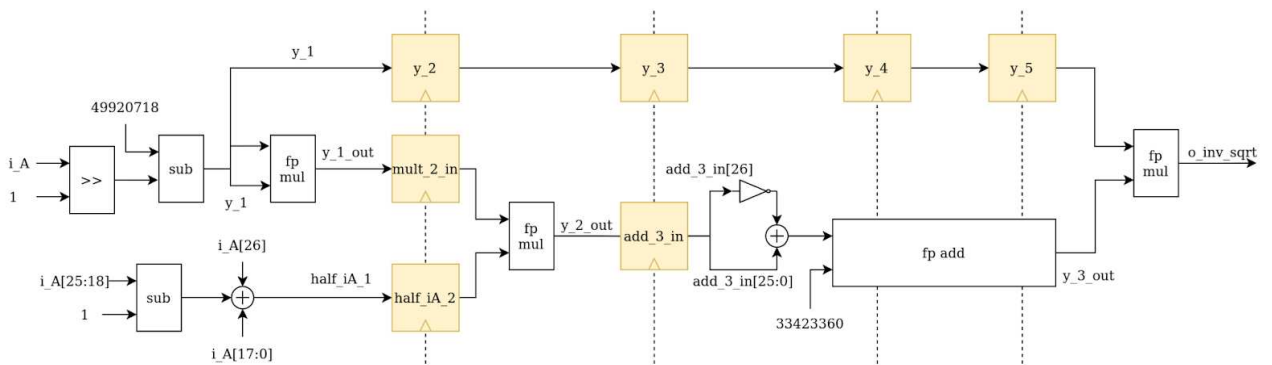


Figure 5: Fp Fast Inverse Square Root Module

#### 4.1.4 Pipelined Two-Body Core

Each two-body core implements the full softened pairwise acceleration computation as a single fully pipelined datapath, accepting a new body pair every clock cycle (5). The pipeline proceeds in five stages, illustrated in Figure 6:

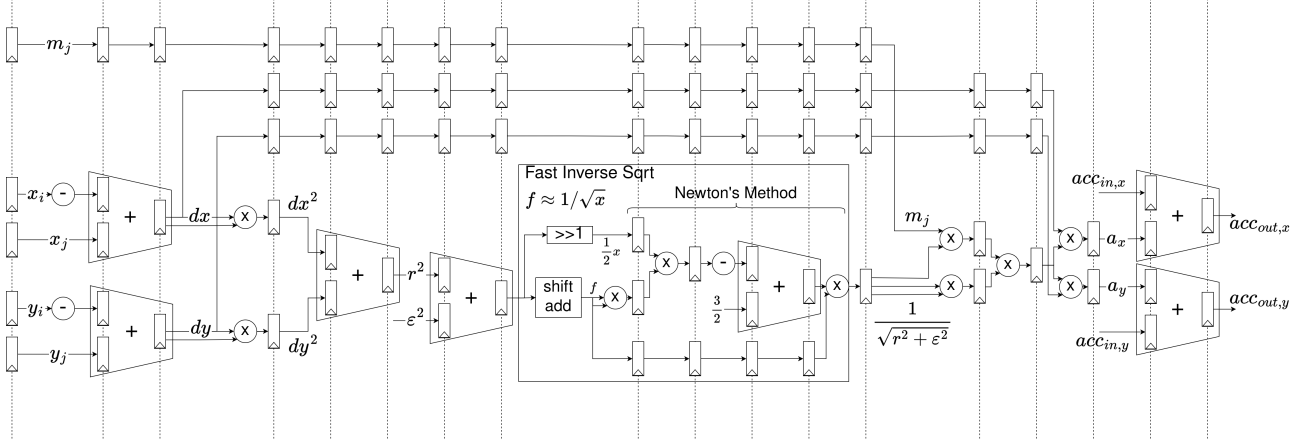


Figure 6: Pipelined datapath of a single two-body compute core. Inputs  $x_i, x_j, y_i, y_j, m_j$  enter from the left; the Fast Inverse Square Root unit occupies the central stages; and the final stages scale and accumulate the pairwise acceleration contributions into `acc_out_x` and `acc_out_y`. Vertical dashed lines denote pipeline stage boundaries; rectangular blocks denote pipeline registers.

1. **Displacement:** Compute  $\Delta x = x_i - x_j$  and  $\Delta y = y_i - y_j$  in parallel using floating-point subtractors.
2. **Squared distance:** Compute  $\Delta x^2, \Delta y^2$ , sum them, and add  $\varepsilon^2$  to obtain the softened  $r^2 + \varepsilon^2$ .
3. **Fast inverse square root:** Apply the FISR pipeline to compute  $f \approx 1/\sqrt{r^2 + \varepsilon^2}$ , then multiply  $f^2$  and apply one Newton–Raphson correction to obtain  $f^3 \approx (r^2 + \varepsilon^2)^{-3/2}$ .
4. **Force scaling:** Multiply  $f^3$  by  $m_j$  to obtain the scalar force coefficient, then multiply separately by  $\Delta x$  and  $\Delta y$  to obtain  $a_x$  and  $a_y$ .
5. **Accumulation:** Add  $a_x$  and  $a_y$  to the running partial sums `acc_in_x` and `acc_in_y`, producing `acc_out_x` and `acc_out_y`.

Four such cores operate in parallel, each assigned a distinct body pair at each clock cycle by the controller.

## 4.2 N-body Memory

The N-body memory module is the main on-chip storage block for the simulation state. It stores each body’s position, velocity, mass, and acceleration fields, with the major attributes placed in separate M10K-backed arrays so that different fields can be accessed and updated independently. During initialization, the HPS software writes the initial body records into this memory through the Avalon interface. During hardware execution, the controller reads body data from memory, the acceleration pipeline writes back accumulated acceleration values, and the leapfrog integrator writes back updated position and velocity values. The memory uses synchronous reads, so requested body data becomes valid one clock cycle after the read address is issued. Internally, write-selection logic resolves updates from the initialization path, acceleration writeback path, and integrator writeback path so that each pipeline stage updates only the fields it owns.

```

(* ramstyle = "M10K" *) logic [DATA_W-1:0] x_mem [0:MAX_BODIES-1];
(* ramstyle = "M10K" *) logic [DATA_W-1:0] y_mem [0:MAX_BODIES-1];
(* ramstyle = "M10K" *) logic [DATA_W-1:0] m_mem [0:MAX_BODIES-1];
(* ramstyle = "M10K" *) logic [DATA_W-1:0] vx_mem [0:MAX_BODIES-1];
(* ramstyle = "M10K" *) logic [DATA_W-1:0] vy_mem [0:MAX_BODIES-1];
(* ramstyle = "M10K" *) logic [DATA_W-1:0] ax_mem [0:MAX_BODIES-1];
(* ramstyle = "M10K" *) logic [DATA_W-1:0] ay_mem [0:MAX_BODIES-1];

```

Figure 7: N-body memory block.

### 4.3 Four-Core Wrapper and Data Flow

The four-core wrapper is the main acceleration datapath used to compute pairwise interactions. Instead of processing one target body at a time, the wrapper works on a tile of 16 target bodies. These 16 bodies are first loaded from N-body memory into a local input bank inside the wrapper. The tile is divided into four groups, where each group contains four target bodies that are mapped onto the four 2-body cores.

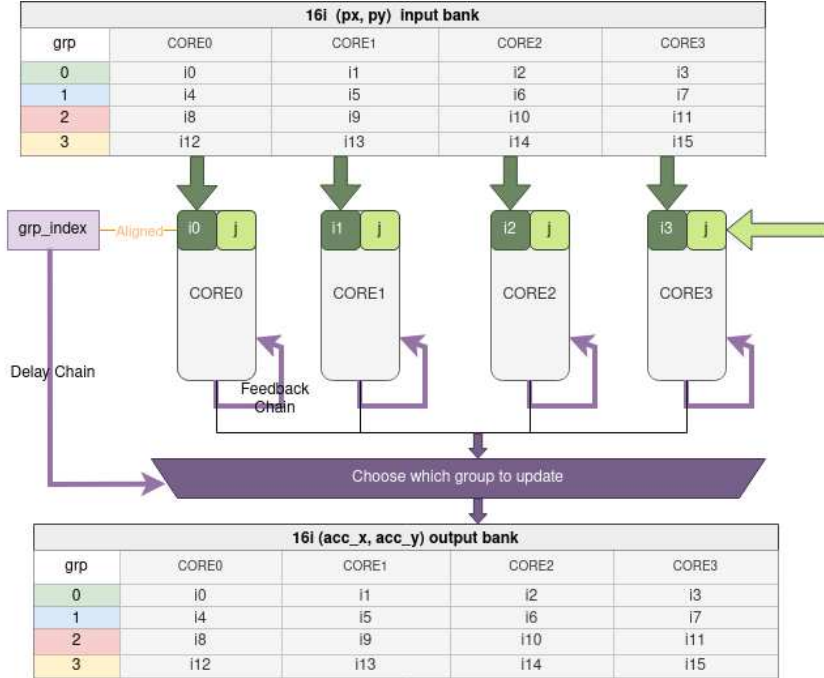


Figure 8: Four-core wrapper dataflow. A 16-body tile is divided into four groups, with each group mapped onto the four 2-body cores.

#### 4.3.1 Tiled Input Organization

The local input bank stores the 16 target bodies as four groups of four bodies. The group index selects which four bodies are currently connected to the four cores. For example, group 0 maps bodies  $i_0$ – $i_3$  to cores 0–3, group 1 maps bodies  $i_4$ – $i_7$ , and so on. This allows the same four physical cores to be reused across all 16 bodies in the tile.

#### 4.3.2 Broadcasted Source Body

During computation, the controller streams one source body  $j$  into the wrapper each cycle. The source body's position  $(x_j, y_j)$  is broadcast to all four cores, while each core receives one target body  $(x_i, y_i)$  from the selected group. The source mass  $m_j$  is also broadcast, but it can be masked per lane. This lane mask is used to disable invalid bodies at the end of a tile and to remove self-interactions where  $i = j$ .

Each 2-body core computes the acceleration contribution from the current source body  $j$  to its assigned target body  $i$ . Because the cores are fully pipelined, a new source body can be issued every cycle while previous interactions continue moving through the pipeline.

### 4.3.3 Accumulation and Feedback

Each 2-body core accumulates acceleration over all source bodies. The previous accumulated acceleration is fed back into the core so that each new pairwise contribution is added to the running sum:

$$a_{i,x}^{new} = a_{i,x}^{prev} + a_{ij,x}, \quad a_{i,y}^{new} = a_{i,y}^{prev} + a_{ij,y}.$$

The feedback path is delayed to align with the pipeline latency of the core. After the pipeline is drained, the accumulated acceleration values represent the total acceleration for the selected group of target bodies.

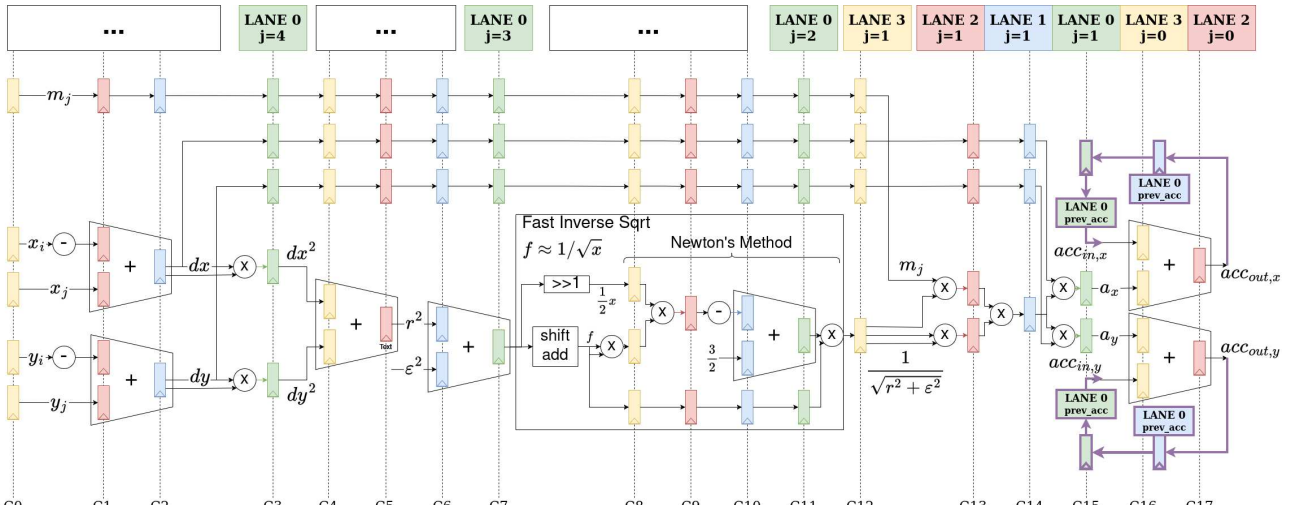


Figure 9: Single 2-body core pipeline. The core computes one pairwise acceleration contribution and feeds the accumulated result back into the pipeline for the next source body.

### 4.3.4 Output Selection and Writeback

The wrapper stores the accumulated acceleration results for all 16 bodies in a local output bank. Since only four cores are available, the controller selects one group at a time and writes the four corresponding acceleration results back to the N-body acceleration memory. The `grp_index` signal is therefore used both to select which input group is currently being computed and to select which output group is being written back. After all four groups are stored, the controller advances to the next 16-body tile or proceeds to the integration stage.

## 4.4 Leapfrog Integrator

### 4.4.1 Leapfrog Integration Algorithm

To update the position and velocity of each body over time, we adopt the leapfrog integration algorithm in hardware, replacing the simpler Euler method (6). Unlike the Euler update:

$$v_{new} = v_{old} + a \cdot dt, \quad pos_{new} = pos_{old} + v_{new} \cdot dt,$$

the leapfrog method splits each timestep into three stages:

$$v_{i+1/2} = v_i + \frac{1}{2}a_i \cdot dt \quad (9)$$

$$x_{i+1} = x_i + v_{i+1/2} \cdot dt \quad (10)$$

$$v_{i+1} = v_{i+1/2} + \frac{1}{2}a_{i+1} \cdot dt \quad (11)$$

This approach uses the mid-interval velocity to update position, yielding better long-term accuracy than Euler integration and preserving energy conservation over extended simulations. Critically, Steps 9, 10, and 11 each require only simple multiply-add operations, while the acceleration computation is handled by the two-body cores. To simplify compute, we combine Step 9 with the next iteration’s Step 11, so that leapfrog integration can be done in 2 steps:

$$v_{i+1/2} = v_{i-1/2} + a_i \cdot dt \quad (12)$$

$$x_{i+1} = x_i + v_{i+1/2} \cdot dt \quad (13)$$

except for the first step where

$$v_{1/2} = v_0 + \frac{1}{2}a_0 \cdot dt \quad (14)$$

#### 4.4.2 Integrator Implementation

The leapfrog integrator is implemented as a small floating-point datapath controlled by a mini FSM. It takes the current body state  $(x, y, v_x, v_y)$  and acceleration  $(a_x, a_y)$ , first updates the velocity, and then uses the updated velocity to update the position:

$$\begin{aligned} v'_x &= v_x + a_x, & v'_y &= v_y + a_y, \\ x' &= x + v'_x, & y' &= y + v'_y. \end{aligned}$$

We do not include an explicit multiplication by  $dt$  in the integrator because the timestep scaling is absorbed into the acceleration and velocity increments used by the hardware, effectively treating each hardware update as one normalized timestep. The datapath uses pipelined floating-point adders, so the control FSM inserts wait states between the velocity update and the position update. After both add stages complete, the integrator asserts `o_done` to notify the main controller that the updated body state is ready.

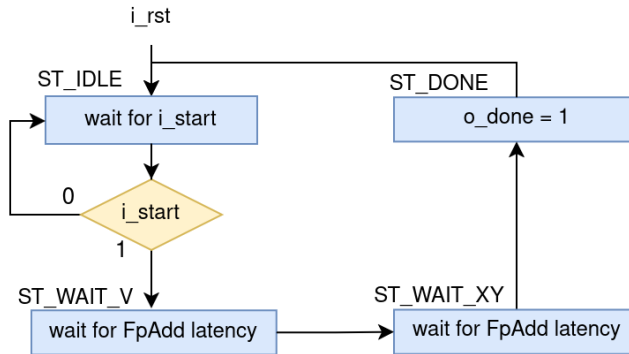


Figure 10: Mini FSM for the leapfrog integrator. The FSM waits for `i_start`, accounts for the floating-point adder latency during the velocity and position update stages, and asserts `o_done` when integration is complete.

## 4.5 N-body Controller

The accelerator controller is implemented as a finite state machine that sequences the full hardware simulation step. At a high level, the controller loads a tile of up to 16 bodies into the four-core wrapper, streams every source body  $j$  through the 2-body cores, stores the accumulated acceleration results, runs the leapfrog integrator for each body, and repeats this process for the requested number of internal timesteps set by `gap`. Figure 11 shows the controller FSM.

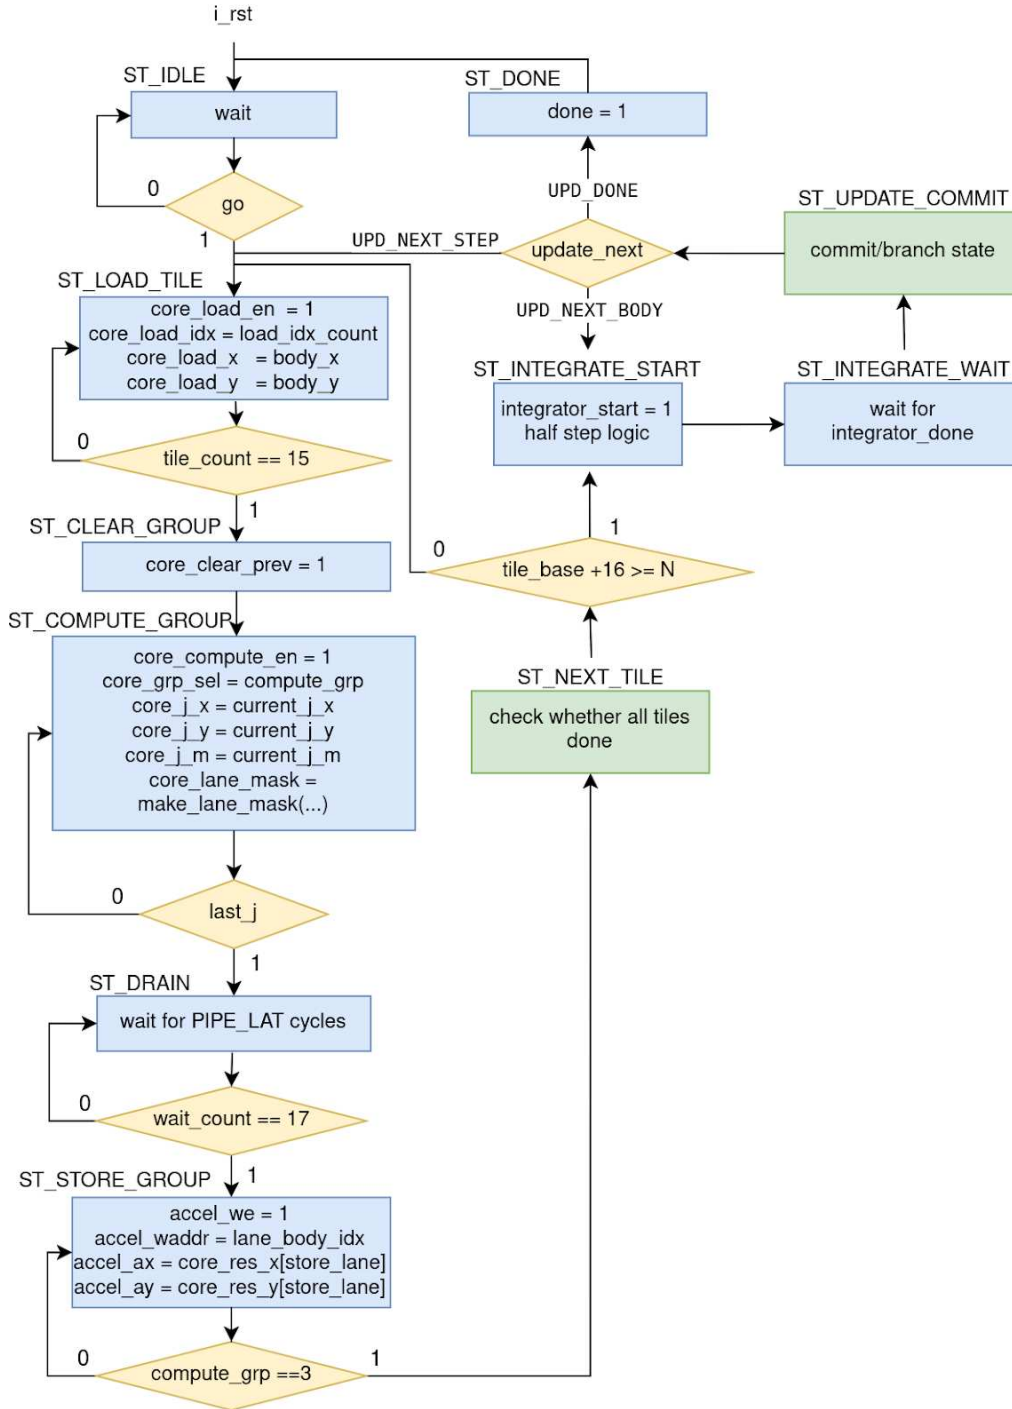


Figure 11: N-body Controller

The `_PRIME` states are one-cycle setup states used before data-dependent operations. They allow the controller to issue a memory read address first, then use the returned body data in the following state. Therefore, `ST_LOAD_TILE_PRIME` is grouped with `ST_LOAD_TILE`, `ST_COMPUTE_PRIME` is grouped with `ST_COMPUTE_GROUP`, and `ST_INTEGRATE_PRIME` is grouped with `ST_INTEGRATE_START`.

The main computation proceeds tile by tile. For each tile, the controller first loads 16 target bodies into the core wrapper. It then broadcasts each source body  $j$  to the four 2-body cores and cycles through four compute groups, allowing the hardware to accumulate acceleration contributions for all 16 target bodies. After the last source body has been processed, the controller waits for the pipeline to drain and stores the four-lane acceleration outputs group by group.

Table 1: Main controller state groups.

State group	Function	Main outputs
ST_IDLE	Wait for software to assert <code>go</code> . Initialize counters, tile index, timestep count, and active body count.	-
ST_LOAD_TILE_PRIME ST_LOAD_TILE	/ Read body data from N-body memory and load one tile of up to 16 target bodies into the four-core wrapper.	<code>body_raddr</code> , <code>core_load_en</code> , <code>core_load_idx</code>
ST_CLEAR_GROUP	Clear previous accumulated acceleration values before computing a new tile.	<code>core_clear_prev</code>
ST_COMPUTE_PRIME ST_COMPUTE_GROUP	/ Stream each source body $j$ through the core wrapper. The controller cycles through four groups and applies lane masks for inactive bodies and self-interactions.	<code>core_compute_en</code> , <code>core_grp_sel</code> , <code>core_lane_mask</code>
ST_DRAIN_GROUP	Wait for the fully pipelined 2-body cores to drain so the final accumulated results are valid.	<code>wait_count</code>
ST_STORE_GROUP ST_NEXT_GROUP	/ Store computed acceleration results from each lane into acceleration memory, then advance to the next compute group.	<code>accel_we</code> , <code>accel_waddr</code> , <code>accel_ax/ay</code>
ST_NEXT_TILE	Check whether all 16-body tiles have been processed. If not, advance to the next tile.	<code>tile_base</code>
ST_INTEGRATE_PRIME ST_INTEGRATE_START ST_INTEGRATE_WAIT	/ Run the leapfrog integrator for each body using the stored acceleration values.	<code>integrator_start</code> , <code>body_raddr</code>
ST_UPDATE_COMMIT	Write updated position and velocity back into N-body memory, then branch to the next body, next timestep, or done state.	<code>body_update_we</code> , <code>body_update_addr</code>
ST_DONE	Assert completion so software can read the updated positions.	<code>done</code>

After all tiles have been processed, the controller enters the integration phase. The integrator is started once per body, and when `integrator_done` is asserted, the updated position and velocity are written back into the N-body memory. The `ST_UPDATE_COMMIT` state then chooses one of three transitions: continue to the next body, start the next internal timestep if `gap` has not been reached, or enter `ST_DONE` when the full update is complete.

## 5 Software Implementation

The software is responsible for system initialization, user interaction, accelerator control, frame history management, and VGA rendering. The expensive all-pairs N-body computation is offloaded to the FPGA accelerator, while the HPS software coordinates when new simulation frames are computed, read back, stored, and displayed.

The software stack is organized into three user-space threads and two Linux kernel drivers. The user-space application communicates with the accelerator through `/dev/nbody` and with the VGA framebuffer through `/dev/nbody_display`. Both device files are implemented as Linux misc devices backed by platform drivers that map the corresponding Avalon-MM peripherals from the device tree.

## 5.1 Software Architecture

At startup, the program takes two command-line arguments: the number of bodies and the GAP value. The number of bodies is limited to `MAX_BODIES = 1024`, while the GAP value is limited to `1--10`. The GAP value controls how many internal hardware timesteps are executed between successive software readbacks, allowing the user to trade off visual update rate and simulation speed.

After argument checking, the main program allocates the frame history buffer, initializes the random seed, resets the simulation state, and creates three threads: the accelerator thread, the display thread, and the USB keyboard thread. The accelerator thread communicates with the N-body hardware, the display thread renders frames into the VGA framebuffer, and the keyboard thread updates shared simulation state based on user input.

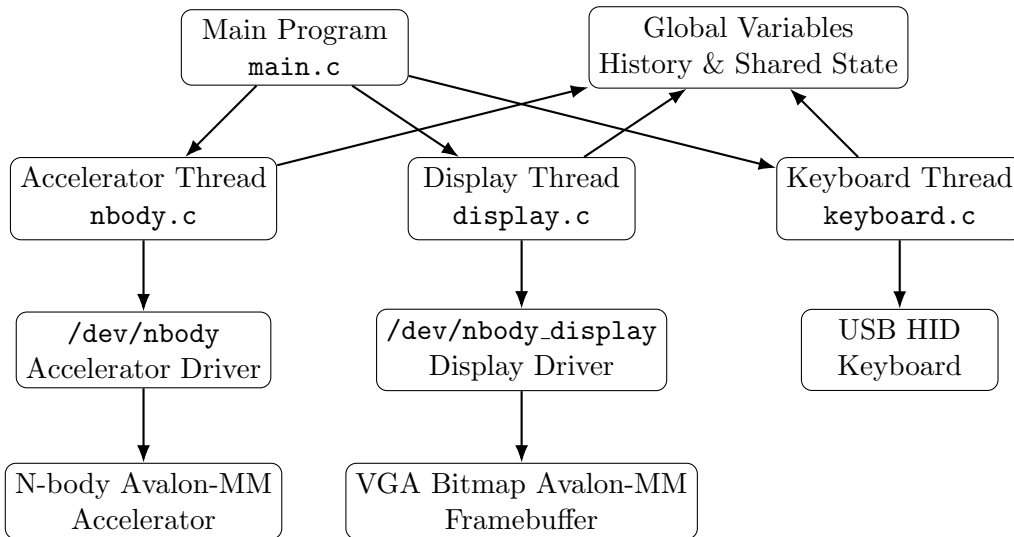


Figure 12: Software architecture. The application uses separate threads for accelerator control, display rendering, and keyboard input. Kernel drivers translate high-level `ioctl` commands into Avalon-MM transactions.

## 5.2 Real-Time Multi-Threaded Software Design

A key software contribution of this project is that the HPS program is structured as a real-time, multi-threaded N-body simulator rather than a simple sequential hardware test program. The application separates accelerator control, display rendering, and user input into three independent threads. This organization is important because these three tasks have very different timing behavior.

The accelerator thread is hardware-latency dependent. It starts a hardware simulation step, polls the accelerator until completion, reads back the updated body positions, and commits the result into the frame history buffer. The display thread is frame-rate dependent. It periodically selects the current history frame, renders it into a packed software framebuffer, and sends the complete frame to the VGA display driver. The keyboard thread is event-driven. It continuously monitors the USB keyboard and updates shared control state such as pause, reset, GAP, and history-frame index.

If these operations were implemented in a single sequential loop, the user interface would become less responsive. For example, keyboard input could only be processed after a hardware run and a full framebuffer update finished. Similarly, display refreshes could stall while the software is waiting for the accelerator to assert `DONE`. By using three threads, the program allows hardware computation, display refresh, and user interaction to progress concurrently. This makes the final system behave like an interactive simulator rather than a batch-mode accelerator demo.

The thread-level division also decouples the hardware computation rate from the display update rate. The accelerator can compute new simulation frames as quickly as the FPGA allows, while the display thread renders at a controlled visual frame rate. The history buffer sits between these two rates: newly computed frames are appended by the accelerator thread, while the display thread can show either the newest frame or a user-selected historical frame. This producer-consumer structure enables smooth playback, pause/resume, reset, and backward/forward frame navigation without requiring the accelerator to recompute old states.

Shared state is protected using `state_mutex`, while hardware `ioctl` access is protected using a separate `hw_mutex`. This separation keeps long hardware transactions from unnecessarily blocking display-state reads and keyboard updates. As a result, the software design supports real-time interaction while still safely coordinating access to shared simulation data and memory-mapped hardware peripherals.

### 5.3 Major Software Components

Table 2: Major software files and responsibilities.

File	Responsibility
<code>main.c</code>	Parses command-line arguments, allocates the history buffer, initializes the random seed, and starts the accelerator, display, and keyboard threads.
<code>nbody.c/h</code>	Maintains global simulation state, initializes random particles, controls pause/reset/history behavior, starts hardware runs, polls for completion, and stores returned body positions.
<code>accelerator_driver.c</code>	Implements the <code>/dev/nbody</code> kernel driver. It maps the accelerator Avalon-MM registers, converts between IEEE single-precision floats and the custom 27-bit floating-point format, and provides <code>ioctl</code> commands for configuration, start, done polling, and result readback.
<code>display.c/h</code>	Converts body positions to screen coordinates, draws particle circles and text into a packed 1-bpp software framebuffer, and sends each frame to the display driver.
<code>display_driver.c</code>	Implements the <code>/dev/nbody_display</code> kernel driver. It copies a packed user-space framebuffer and writes all 9600 32-bit words into the VGA Avalon-MM framebuffer.
<code>keyboard.c/h</code> <code>usbkeyboard./h</code>	Find a USB HID keyboard using <code>libusb</code> , decode keycodes, and call the corresponding simulation-control functions.

### 5.4 Shared Simulation State and History Buffer

The main shared software state is defined in `nbody.c`. It includes the active body count, current GAP value, pause state, running state, static display radius for each body, and a history buffer storing previously computed frames.

Each history entry stores only body positions:

$$\text{body\_pos\_t} = (x, y),$$

where both coordinates are single-precision floating-point values in software. The history buffer contains `MAX_HISTORY = 32768` frame slots, and each slot can store up to `MAX_BODIES = 1024` body

positions. Therefore, the maximum allocated history storage is

$$32768 \times 1024 \times 2 \times 4 = 268,435,456 \text{ bytes} \approx 256 \text{ MiB.}$$

This relatively large buffer allows the user to scrub through many previously computed frames without recomputing them in hardware.

The history buffer is used as a ring buffer. The variable `h_head` points to the next history slot to be written, `h_count` records how many valid frames are available, and `view_idx` selects the frame currently shown by the display thread. Since `MAX_HISTORY` is a power of two, wrap-around is implemented efficiently using

$$\text{slot} = \text{h\_head} \& (\text{MAX\_HISTORY} - 1).$$

This avoids the need for expensive modulo operations in the frame-commit path.

Two mutexes are used to protect shared resources. The `state_mutex` protects simulation state variables such as pause state, history counters, current view index, and radius indices. The `hw_mutex` protects hardware access through `ioctl` calls. Separating these locks prevents display and keyboard operations from being blocked unnecessarily by longer hardware transactions.

## 5.5 Initialization and Reset Flow

The initialization and reset paths share the same software routine, `reset_system()`. During reset, software randomly generates a new set of initial particles. The position range is

$$x, y \in [-10, 10],$$

the mass range is

$$m \in [0.002, 0.01],$$

and the initial velocities are set to

$$v_x = v_y = 0.$$

The initial body positions are written into `history[0]` so that the display can show the first frame immediately after reset. The body mass is also mapped into one of four display radius indices, so heavier bodies are drawn with larger circles.

After random initialization, the software writes the body count and GAP configuration to the accelerator driver using `NBODY_WRITE_CONFIG`. It then writes all body records using `NBODY_WRITE_BODIES`. The kernel driver converts each software `float` into the accelerator's custom S1E8M18 27-bit floating-point format before writing the values to the Avalon-MM input registers.

A generation counter, `sim_generation`, is used to prevent stale hardware results from being stored after a reset. If the user resets the system while the accelerator is still computing an old frame, the returned old result is discarded rather than inserted into the new simulation's history buffer. This makes reset behavior reliable even though the accelerator thread and keyboard thread run concurrently.

## 5.6 Accelerator Thread

The accelerator thread controls the hardware simulation loop. It opens `/dev/nbody`, loads the initial simulation state, and repeatedly starts hardware computation when the simulation is not paused. For each frame, the thread performs the following sequence:

1. Select the next available history-buffer slot.

2. Issue `NBODY_START_RUN`, which causes the driver to pulse the accelerator `GO` register.
3. Poll `NBODY_CHECK_DONE` until the hardware asserts `DONE = 1`.
4. Read all updated body positions using `NBODY_READ_RESULTS`, asserting `READ = 1` and reading `OUT_X/OUT_Y` for each body.
5. Clear the hardware read state using `NBODY_CLEAR_READ`, lowering `READ = 0`.
6. If no reset occurred during this run, commit the frame into the history ring and update `view_idx`; otherwise discard the result.

Figure 13 illustrates this control flow, including the hardware signals exchanged at each step.

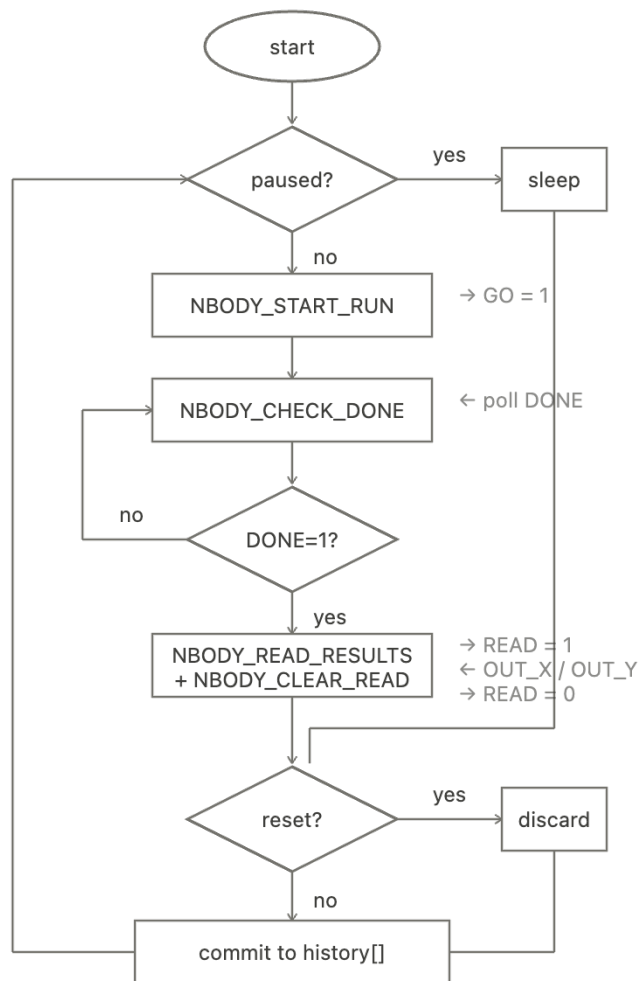


Figure 13: Accelerator thread control flow

The high-level accelerator loop is summarized below.

Listing 1: High-level accelerator-thread control loop.

```

while (running) {
    if (is_paused) {
        sleep briefly;
        continue;
    }
}
  
```

```

next_idx = next history slot;

ioctl(nbody_fd, NBODY_START_RUN);

do {
    ioctl(nbody_fd, NBODY_CHECK_DONE, &done);
    sleep briefly if not done;
} while (running && !done);

ioctl(nbody_fd, NBODY_READ_RESULTS, history[next_idx]);
ioctl(nbody_fd, NBODY_CLEAR_READ);

if (no reset occurred during this run) {
    commit next_idx into the history ring;
    view_idx = newest frame;
}
}

```

The design uses polling rather than interrupt-driven completion. This keeps the driver simple and is sufficient for the current interactive application, because the polling loop sleeps briefly between DONE checks and the display runs in a separate thread.

## 5.7 Display Thread and Frame Rendering

The display thread is responsible for converting simulation frames into a  $640 \times 480$  packed 1-bit-per-pixel framebuffer. The top 448 pixel rows are used for body visualization, while the bottom 32 pixels are reserved for the text-based user interface. This matches two rows of  $8 \times 16$  font characters.

Before rendering, the display thread copies the currently selected history frame into a local temporary buffer while holding `state_mutex`. The mutex is then released before drawing, so the accelerator and keyboard threads are not blocked during framebuffer generation.

World coordinates are mapped to screen coordinates using linear normalization:

$$x_{\text{screen}} = \frac{x - x_{\min}}{x_{\max} - x_{\min}}(640 - 1),$$

$$y_{\text{screen}} = \frac{y - y_{\min}}{y_{\max} - y_{\min}}(448 - 1),$$

where  $x_{\min} = y_{\min} = -10$  and  $x_{\max} = y_{\max} = 10$ . This mapping keeps the simulated world coordinate range aligned with the visible display area.

### 5.7.1 Precomputed Circular Body Masks

A small but important rendering optimization is the use of precomputed circular body masks. The display code supports four particle radii,

$$r \in \{1, 2, 3, 4\},$$

which correspond to four mass ranges. Heavier bodies are assigned larger radius indices during initialization, so the display gives a simple visual indication of mass without requiring color or grayscale output.

Instead of recomputing the circle equation for every body in every frame, the display thread precomputes all circular masks once during display initialization. These masks are stored in a lookup table:

```
body_masks [NUM_BODY_RADII] [9] [9].
```

The largest supported body radius is 4, so a  $9 \times 9$  window is sufficient to cover all pixel offsets from  $-4$  to  $+4$  in both the horizontal and vertical directions.

For each supported radius, the initialization routine scans all local offsets  $(x, y)$  around the particle center and marks a mask entry as active if the point lies inside the filled circle:

$$x^2 + y^2 < r^2 + r.$$

The additional  $+r$  term slightly thickens the discrete circle boundary, producing more visually solid particles on a low-resolution 1-bpp display. This is especially useful for small radii, where a strict  $x^2 + y^2 < r^2$  test can produce sparse or uneven-looking shapes.

During frame rendering, the function that draws a body no longer evaluates the circle equation. For each particle, the renderer selects the precomputed mask using the particle's `radius_idx`, loops over the local  $9 \times 9$  mask window, and calls `set_pixel()` only for mask entries that are set. This changes the per-particle drawing operation from “compute geometry and draw” to “lookup mask and draw,” which is simpler and more predictable inside the real-time display loop.

Listing 2: Precomputed circular body mask generation and reuse.

```
/* Initialization: done once */
for each radius r in {1, 2, 3, 4}:
    for y = -r to r:
        for x = -r to r:
            if (x*x + y*y < r*r + r)
                body_mask[r][y+r][x+r] = 1;

/* Rendering: reused for every body in every frame */
for each body:
    select mask using radius_idx;
    for each local pixel in the mask:
        if mask pixel is set:
            set_pixel(center_x + x, center_y + y);
```

This optimization is especially useful because the display thread may need to draw up to 1024 particles every frame. Since the body shapes are static and only the particle centers change over time, precomputing the masks removes repeated arithmetic from the frame loop while preserving the appearance of circular particles. It also keeps the display code independent from the accelerator: the FPGA produces updated physical positions, while the HPS software handles the final visualization step efficiently in user space.

The software framebuffer uses the same packed format as the hardware display peripheral:

$$\text{word\_index} = y \cdot 20 + \left\lfloor \frac{x}{32} \right\rfloor, \quad \text{bit\_index} = x \bmod 32.$$

A set bit corresponds to a white pixel, while a cleared bit corresponds to a black pixel. After all particles and UI text are drawn, the thread calls `DISPLAY_WRITE_FRAME` to send the complete 9600-word framebuffer to `/dev/nbody_display`. The display thread sleeps for approximately 33.3 ms between frames, giving an intended software rendering rate of about 30 frames/s.

## 5.8 Keyboard Input and User Interface

The application is controlled entirely through a USB keyboard. The keyboard thread uses `libusb` to locate a USB HID keyboard and reads interrupt-transfer packets from the keyboard endpoint. USB HID keycodes are then translated into simulation actions.

Table 3: Keyboard controls for the simulation.

Key	Action
SPACE	Pause or resume the simulation
R	Reset the simulation with newly randomized bodies
W	Increase GAP, causing more hardware timesteps per displayed readback
S	Decrease GAP, causing fewer hardware timesteps per displayed readback
A	Step backward by one stored history frame
D	Step forward by one stored history frame
Q	Quit the application

Manual frame navigation automatically pauses the simulation so that the user can inspect the stored history. The A and D keys also support long-press repeat: after an initial delay, the keyboard thread repeatedly steps through history at a fixed interval while the key is held. Other keys are de-bounced by comparing the current USB packet against the previous packet, avoiding duplicate actions from a single key press.

The display thread also renders a simple text-based UI in the bottom region of the VGA screen. This status line shows the number of bodies, the current GAP value, the current history-frame index, the latest computed frame, and whether the simulation is running or paused. This makes the FPGA demo self-contained: the user can adjust simulation behavior and immediately see both the visual result and the current software state on the VGA display.

## 5.9 Kernel Driver Layer

The software uses two kernel drivers to isolate low-level Avalon-MM access from the user-space application.

The accelerator driver registers the device file `/dev/nbody`. It obtains the accelerator's physical address range from the device tree, maps the Avalon-MM register space using `of_iomap()`, and exposes high-level `ioctl` commands for software control. The driver is also responsible for converting data between user-space IEEE-754 single-precision floats and the accelerator's custom 27-bit S1E8M18 format. On input, the 23-bit IEEE mantissa is rounded down to 18 mantissa bits and stored in the lower 27 bits of a 32-bit Avalon word. On output, the 27-bit hardware result is expanded back into a 32-bit software float.

The display driver registers the device file `/dev/nbody_display`. It maps the VGA framebuffer peripheral, allocates a kernel-side framebuffer buffer, and provides two commands: `DISPLAY_WRITE_FRAME` and `DISPLAY_CLEAR`. For each frame update, it copies the 9600-word packed framebuffer from user space and writes each word into the hardware framebuffer using `iowrite32`. The driver does not perform any rendering; all particle drawing and text drawing are handled in user space.

This driver structure gives the user-space application a clean interface to both FPGA peripherals. Instead of directly manipulating physical addresses, the application sends structured `ioctl` requests. The kernel drivers then perform the required memory-mapped I/O operations in the correct order.

## 5.10 End-to-End Software Flow

Overall, each interactive simulation frame follows this end-to-end flow:

1. User space initializes or resets body positions, masses, and velocities.
2. The accelerator driver converts the body data to 27-bit floating-point format and writes the data into the hardware input registers.
3. The accelerator thread starts a hardware run and polls until `DONE` is asserted.
4. The accelerator driver reads updated positions from the hardware output stream and converts them back to software floats.
5. The accelerator thread stores the returned positions into the history ring buffer.
6. The display thread selects either the newest frame or a user-selected historical frame.
7. The display thread renders particles and UI text into the packed software framebuffer.
8. The display driver copies the packed framebuffer into the VGA peripheral for display.

This structure separates the system into clear responsibilities: the FPGA performs the numerical N-body update, the kernel drivers provide safe and structured access to Avalon-MM peripherals, and the user-space application handles control flow, history management, rendering, and interaction. Together, these components form a complete real-time hardware/software N-body simulator rather than only a standalone accelerator benchmark.

## 6 Hardware/Software Interface

### 6.1 Accelerator Interface

The accelerator is controlled through a memory-mapped Avalon interface. Each transaction transfers one 32-bit word; for 27-bit floating-point values, bits [26:0] carry the data and bits [31:27] are ignored on write and return zero on read.

Body data are written sequentially through dedicated input registers. For each body, software must write the five input fields in the fixed order `X_IN`, `Y_IN`, `M_IN`, `VX_IN`, `VY_IN`. Writing `VY_IN` is treated as the completion of the current body entry; the hardware commits the assembled body data into internal memory and automatically increments the internal input pointer to the next body.

Output readback is symmetric with input streaming. After `DONE` is asserted, software reads `OUT_X` followed by `OUT_Y` for each body in sequence. Reading `OUT_Y` increments the internal output pointer to the next body. Both the input and output pointers are implicitly reset to zero when `GO` is asserted.

### 6.2 Display Interface

The display is controlled through a separate memory-mapped Avalon peripheral, `vga_bitmap_avmm`. Both the accelerator and display peripherals are mapped into the HPS lightweight Avalon bridge address space, whose physical base address is `0xFF200000`. To avoid address conflicts, the accelerator peripheral starts at offset `0x00000000`, while the display peripheral starts at offset `0x00010000`. Therefore, software accesses the display frame buffer using addresses relative to the display peripheral base.

Address	Register	R/W	Description
0x00	GO	W	Pulse high to start computation. Implicitly resets both input and output body pointers to 0.
0x01	N_BODIES	W	Number of active bodies in the simulation.
0x02	GAP	W	Number of timesteps executed internally between successive DONE assertions.
0x03	X_IN	W	Input X position for the current body.
0x04	Y_IN	W	Input Y position for the current body.
0x05	M_IN	W	Input mass for the current body.
0x06	VX_IN	W	Input X velocity for the current body.
0x07	VY_IN	W	Input Y velocity for the current body.
0x08	DONE	R	Set high by hardware upon completion of GAP timesteps. Cleared when software lowers READ.
0x09	READ	W	Asserted high by software after DONE is observed. Lowered by software once all outputs have been read, signalling readiness for the next GO.
0x0A	OUT_X	R	Output X position for the current body.
0x0B	OUT_Y	R	Output Y position for the current body.

The display register map is organized as a packed 1-bit-per-pixel frame buffer instead of a set of control registers. Software renders each frame in user space, packs the  $640 \times 480$  bitmap into 32-bit words, and sends the complete frame to the kernel driver through `/dev/nbody_display`. The display driver then writes the packed frame buffer to the Avalon address space of the VGA peripheral.

Table 4: Display peripheral memory map.

Address Range	Name	R/W	Description
0x0000-0x257F	FRAMEBUFFER	W	Packed $640 \times 480$ 1-bpp frame buffer. Each 32-bit word stores 32 horizontal pixels.

The frame buffer contains 9600 32-bit words, corresponding to

$$640 \times 480 / 32 = 9600$$

words per frame. Since each row has 640 pixels, each row occupies 20 words. Therefore, the word address and bit index for a pixel at screen coordinate  $(x, y)$  are:

$$\text{word\_index} = y \cdot 20 + \left\lfloor \frac{x}{32} \right\rfloor, \quad \text{bit\_index} = x \bmod 32.$$

A bit value of 1 displays a white pixel, while a bit value of 0 displays a black pixel.

The hardware VGA controller continuously scans the frame buffer while software updates it through the Avalon interface. The controller generates the  $640 \times 480$  VGA timing signals and uses the current pixel coordinate to read the corresponding frame-buffer word. The selected bit is converted into RGB output: white for active body pixels and black for the background. Since the frame buffer is implemented as a dual-port memory, software writes and VGA scanout can occur independently.

The kernel driver provides two main operations through `/dev/nbody_display`. The `DISPLAY_WRITE_FRAME` command copies a full packed frame from user space and writes all 9600 words into the hardware frame buffer. The `DISPLAY_CLEAR` command writes zero to every frame-buffer word, clearing the display.

## 7 Testbench and Verification

### 7.1 Arithmetic Unit Testbenches

Before verifying the full N-body datapath, we first designed unit-level SystemVerilog testbenches for the major arithmetic and memory modules. The floating-point arithmetic tests were driven by Python-generated test vectors, so that the RTL outputs could be compared against expected software-computed results.

#### 7.1.1 Floating Pint Adder

The floating-point adder testbench verifies the correctness of the two-stage custom 27-bit floating-point adder in S1E8M18 format. The testbench reads Python-generated test vectors from `fpadd_cases.txt`. Each vector contains two 27-bit input operands and one expected 27-bit output.

Since `FpAdd` is a two-stage pipelined module, the testbench applies each input pair, waits for two positive clock edges, and then compares `oSum` with the expected Python output. The test cases include edge cases like addition with zero, cancellation between opposite inputs, different-sign addition, commutativity checks, and randomized vectors. The adder passed all the directed corner-case tests and randomized Python-generated test vectors.

#### 7.1.2 Floating Pint Multiplier

The floating-point multiplier testbench verifies the combinational 27-bit custom floating-point multiplier using the `fpmul_cases.txt`. The multiplier computes the output sign using the XOR of the input signs, adds the two exponents, and multiplies the mantissa fields to produce final results.

The testbench covers multiplication by zero, positive and negative input combinations, multiplication commutativity, maximum-value multiplication, and randomized input vectors. For each test case, the RTL result was compared against the Python-generated expected output, and the multiplier passed all directed and randomized tests.

#### 7.1.3 Fast Inverse Square Root

The fast inverse square root testbench verifies the pipelined module used to approximate  $1/\sqrt{x}$ . Since the module is pipelined, the testbench also checks output timing and throughput behavior, not only numerical correctness. Our testbench first allows the pipeline to settle and then holds each input constant for `SETTLE_CYCLES = 8` cycles before checking the output. The test cases include accuracy checks, output stability checks, pipeline timing checks, and randomized inputs. For each test case, the RTL output is compared directly against the expected output, and the module passed the Python-generated test cases under the configured pipeline settling delay.

#### 7.1.4 N-body Memory

The N-body memory testbench verifies memory write behavior, read behavior, address independence, field preservation, and write-priority handling of the on-chip memory module that stores the body state. It also verifies the operational behavior such as correct body-data storage, acceleration updates, integrator writeback behavior, independent address access, and simultaneous write-priority

handling. The testbench uses task-based reads and writes to confirm that body updates correctly modify  $x/y/vx/vy$ , acceleration updates only modify  $ax/ay$ , and mass values are preserved across different writeback operations. Additional tests verify synchronous read behavior and simultaneous multi-source writes to the same address. The memory module passed all read/write and write-priority verification tests.

## 7.2 Golden Model

The verification framework uses a Python-based golden model as the software reference for validating the FPGA accelerator outputs. The model emulates the custom FP27 datapath used in the RTL, including the FP27 adder, multiplier, and Fast Inverse Square Root unit. It also follows the same N-body acceleration computation flow and arithmetic ordering as the hardware, allowing the generated acceleration outputs to be directly compared against RTL simulation and FPGA execution results.

To evaluate the numerical accuracy of the custom 27-bit S1E8M18 floating-point format, the accelerator/golden-model outputs are compared against a double-precision Float64 software reference using the same initial frame data. The relative acceleration error is measured across multiple simulation frames.

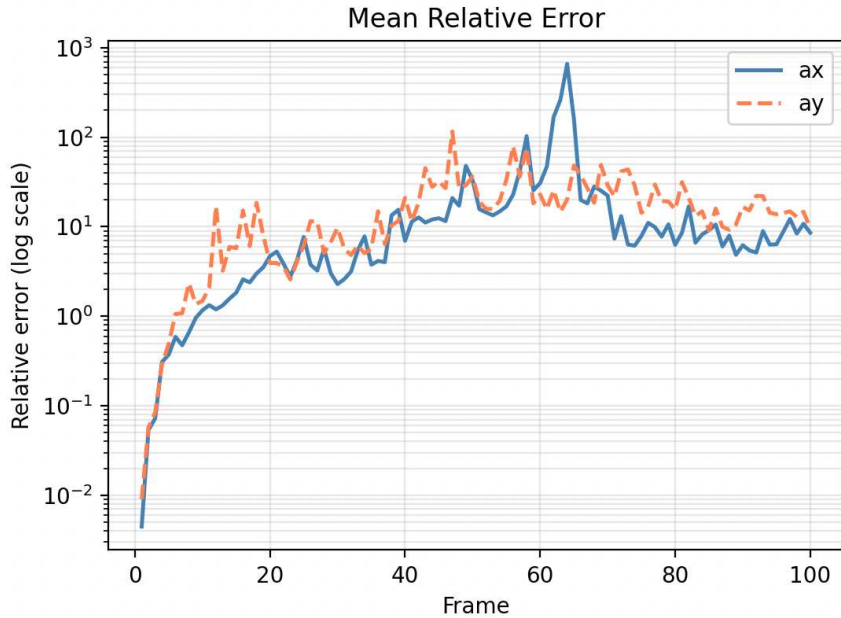


Figure 14: Accuracy validation of the 27-bit S1E8M18 floating-point format against Float64 reference computation.

Figure 14 shows the mean relative error of the computed acceleration components ( $a_x$  and  $a_y$ ) over 100 simulation frames. The first frame captures the intrinsic precision loss introduced by the reduced-precision S1E8M18 format, while the increasing error in later frames is mainly due to the chaotic nature of the N-body system rather than hardware arithmetic mismatch. In the first frame, the custom 27-bit floating-point representation achieves less than 1% mean relative error compared with the Float64 reference, showing that the reduced-precision datapath is sufficiently accurate for this gravitational simulation workload.

### 7.3 Frame Data Generation

A Python-based frame data generator was developed to produce simulation input datasets for both RTL verification and FPGA testing. The generator creates large-scale particle initialization files containing position, velocity, and mass information for the N-body system using the custom 27-bit floating-point representation adopted in the hardware design.

The generated datasets were used as standardized test inputs throughout the verification flow, enabling consistent evaluation across the golden model, RTL simulation, and FPGA execution. In particular, initialization files such as `frame0_1024binit200_27bits.txt` were used to validate accelerator functionality under large-scale 1024-body simulation workloads.

The frame generator also automated the creation of multiple randomized test cases, allowing extensive functional testing and numerical validation of the hardware datapath under different particle distributions and initial conditions.

### 7.4 Outputs and Comparison

We designed a set of SystemVerilog testbenches to validate all stages of the computation pipeline against the Python-based golden models. The same generated 1024-body frame-data inputs were used across RTL simulations, controller verification, FPGA-oriented tests, and golden-model execution, allowing direct output comparison throughout the accelerator datapath and control flow.

#### 7.4.1 Acceleration Accumulation Outputs

Using the generated 1024-body frame datasets as simulation inputs, we first verified the standalone pipelined 2-body core against the Python golden model. The golden model reproduced the same custom FP27 arithmetic flow used in hardware, including the FP27 adder, multiplier, and fast inverse square root modules. After validating the single-core acceleration results, we integrated four parallel 2-body cores into the wrapped accelerator datapath and repeated the comparison using the same frame data inputs. The accumulated acceleration outputs (ax, ay) from both the standalone core and the four-core wrapped design matched the golden-model outputs, confirming that the parallel accumulation logic, j-body broadcast scheduling, accumulation feedback chain, and multi-core datapath integration preserved numerical correctness across the full compute pipeline.

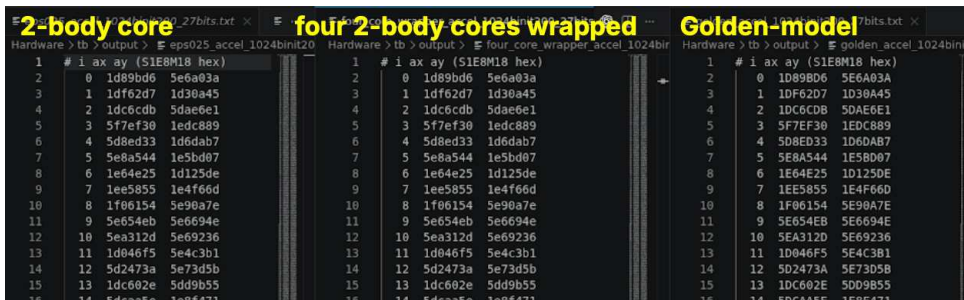


Figure 15: Acceleration Accumulation Outputs comparison

#### 7.4.2 Controller-Integrated Outputs

The controller simulation included tile loading, j-body streaming, acceleration accumulation writeback, leapfrog integration, and final body-state updates. The RTL `nbody_control` module was compared

directly against the Python `golden_control.py` reference model, which emulates the full scheduling and state-update behavior of the hardware pipeline. The final outputs, including particle position ( $x, y$ ), velocity ( $vx, vy$ ), and accumulated acceleration ( $ax, ay$ ), matched the golden-model results across the tested frame data. This verified the correctness of the controller FSM sequencing, memory addressing, write-enable control, integration timing, and overall system-level coordination of the N-body accelerator pipeline.



Figure 16: Controller-Integrated Outputs comparison

### 7.4.3 Software-visible Outputs

We also verified the software-visible outputs through Avalon-MM interface. The `avmm_frame_xy_dump.c` test program loads particle initialization data into the FPGA accelerator through the memory-mapped registers, starts the computation by asserting GO, waits for the DONE signal, and then reads back the final OUT\_X and OUT\_Y values generated by the hardware. These outputs were compared directly against the Python reference generated by `golden_avmm_xy.py`, which models the same software-visible behavior of the accelerator interface. Using the same generated frame-data inputs, the RTL AVMM interface produced the same final body-position outputs as the Python golden model, verifying the correctness of the hardware/ software interface, Avalon-MM control flow, and final output readback path.



Figure 17: Avmm Frame Outputs comparison

## 8 Performance and Resource Utilization

### 8.1 Performance Metrics

The main performance metric for the N-body accelerator is the pairwise interaction throughput, defined as the number of `two_body` interactions computed per second. Since the accelerator instantiates four parallel `two_body` compute units, the theoretical interaction throughput is

$$R_{\text{int}} = 4f_{\text{clk}}.$$

After fitting in Quartus, the worst-case reported FMAX of the full system is

$$f_{\text{max}} = 73.76 \text{ MHz}.$$

This value comes from the slow 1100 mV, 85°C timing model FMAX summary, shown in Figure 18.

```

+-----+
; Slow 1100mV 85C Model Fmax Summary
+-----+-----+-----+
; Fmax      ; Restricted Fmax ; Clock Name
+-----+-----+-----+
; 73.76 MHz ; 73.76 MHz      ; clock_50_1
; 1184.83 MHz ; 717.36 MHz    ; soc_system:soc_system0|soc_system_t
+-----+-----+-----+
This panel reports FMAX for every clock in the design, regardless of

```

Figure 18: Quartus post-fit FMAX summary under the slow 1100 mV, 85°C timing model. The worst reported FMAX is 73.76 MHz for the `clock_50_1` domain.

Therefore, the peak theoretical pairwise interaction throughput is

$$R_{\text{int}} = 4 \times 73.76 \times 10^6 = 295.04 \times 10^6$$

two-body interactions per second.

For  $N = 1024$  bodies, one all-pairs simulation frame requires approximately

$$N^2 = 1024^2 = 1,048,576$$

pairwise interactions. Since four interactions are computed per cycle, the frame latency is

$$\frac{N^2}{4} = \frac{1,048,576}{4} = 262,144$$

accelerator cycles. The theoretical compute frame rate is therefore

$$\text{FPS}_{\text{compute}} = \frac{f_{\text{clk}}}{262,144}.$$

At 73.76 MHz, this gives

$$\text{FPS}_{\text{compute}} = \frac{73.76 \times 10^6}{262,144} \approx 281.4 \text{ frames/s}.$$

The VGA display refresh rate is approximately 60 FPS. Thus, at the post-fit FMAX, the accelerator can compute

$$\frac{281.4}{60} \approx 4.69$$

simulation frames per displayed VGA frame. Equivalently, the display can show approximately every fourth computed frame and skip about three intermediate frames while maintaining 60 FPS. Therefore, for  $N = 1024$ , the accelerator datapath is theoretically fast enough for real-time 60 FPS visualization, assuming that software overhead, Avalon-MM transfer latency, and framebuffer update time do not dominate.

## 8.2 Post-Fit FPGA Resource Utilization

Figure 19 shows the post-fit resource utilization reported by Quartus for the complete DE1-SoC system, including the HPS-connected Avalon-MM peripherals, VGA display logic, memory-mapped accelerator interface, and N-body compute datapath.

```
1  Fitter Status : Successful - Tue May 12 15:48:18 2026
2  Quartus Prime Version : 21.1.0 Build 842 10/21/2021 SJ Lite Edition
3  Revision Name : soc_system
4  Top-level Entity Name : soc_system_top
5  Family : Cyclone V
6  Device : 5CSEMA5F31C6
7  Timing Models : Final
8  Logic utilization (in ALMs) : 14,860 / 32,070 ( 46 % )
9  Total registers : 13576
10 Total pins : 362 / 457 ( 79 % )
11 Total virtual pins : 0
12 Total block memory bits : 499,820 / 4,065,280 ( 12 % )
13 Total RAM Blocks : 60 / 397 ( 15 % )
14 Total DSP Blocks : 40 / 87 ( 46 % )
15 Total HSSI RX PCSs : 0
16 Total HSSI PMA RX Deserializers : 0
17 Total HSSI TX PCSs : 0
18 Total HSSI PMA TX Serializers : 0
19 Total PLLs : 0 / 6 ( 0 % )
20 Total DLLs : 1 / 4 ( 25 % )
```

Figure 19: Quartus post-fit resource utilization summary for the complete system.

The final fitted design uses 14,860 out of 32,070 ALMs, corresponding to 46% of the available logic resources. It also uses 13,576 registers and 362 out of 457 pins. The design uses 499,820 block memory bits out of 4,065,280, corresponding to 12% of the available embedded memory, and 60 out of 397 RAM blocks, corresponding to 15% of the available RAM blocks. The final design uses 40 out of 87 DSP blocks, corresponding to 46% DSP utilization.

Table 5: Post-fit resource utilization of the complete system.

Resource	Used	Available	Utilization
ALMs	14,860	32,070	46%
Registers	13,576	–	–
Pins	362	457	79%
Block memory bits	499,820	4,065,280	12%
RAM blocks	60	397	15%
DSP blocks	40	87	46%
PLLs	0	6	0%
DLLs	1	4	25%

The design uses less than half of the available ALM and DSP resources, and only a small fraction of the available embedded memory. Therefore, the current implementation is not limited by memory capacity. The main implementation constraints are arithmetic resource usage, routing complexity, and timing closure.

### 8.2.1 Custom 27-bit Floating-Point Format

The choice of floating-point precision directly determines how efficiently the design maps onto the Cyclone V DSP blocks (7). The Cyclone V 5CSEA5 FPGA provides 174 independent  $18 \times 18$  bit multipliers in its DSP blocks. The available multiplier configurations and device resource counts are shown in Figure 20.

Variant	Member Code	Variable-precision DSP Block	Independent Input and Output Multiplications Operator			18 x 18 Multiplier Adder Mode	18 x 18 Multiplier Adder Summed with 36 bit Input
			9 x 9 Multiplier	18 x 18 Multiplier	27 x 27 Multiplier		
Cyclone V E	A2	25	75	50	25	25	25
	A4	66	198	132	66	66	66
	A5	150	450	300	150	150	150
	A7	156	468	312	156	156	156
	A9	342	1,026	684	342	342	342
Cyclone V GX	C3	57	171	114	57	57	57
	C4	70	210	140	70	70	70
	C5	150	450	300	150	150	150
	C7	156	468	312	156	156	156
	C9	342	1,026	684	342	342	342
Cyclone V GT	D5	150	450	300	150	150	150
	D7	156	468	312	156	156	156
	D9	342	1,026	684	342	342	342
Cyclone V SE	A2	36	108	72	36	36	36
	A4	84	252	168	84	84	84
	A5	87	261	174	87	87	87
	A6	112	336	224	112	112	112

Figure 20: dsp.blocks.table

The IEEE 754 single-precision floating point uses a 23-bit stored mantissa (24 bits including the implicit leading 1). A  $24 \times 24$  mantissa multiplication exceeds the native  $18 \times 18$  DSP width and must be decomposed using the Karatsuba method:

$$24 = 18 + 6 \implies \text{requires 4 DSP blocks per multiplier}$$

since the four partial products ( $18 \times 18$ ,  $18 \times 6$ ,  $6 \times 18$ ,  $6 \times 6$ ) each occupy one DSP block.

Our custom 27-bit format uses 1 sign bit, 8 exponent bits, and an 18-bit mantissa. An  $18 \times 18$  mantissa multiplication maps exactly onto a single Cyclone V DSP block with no decomposition required. This is the key motivation for the 27-bit format: it is the largest mantissa width that fits within one native DSP block, maximizing precision per DSP block used.

The resulting multiplier density comparison is:

Format	DSP blocks per multiplier	Max parallel multipliers
IEEE 754 32-bit	4	$\lfloor 174/4 \rfloor = 43$
Custom 27-bit	1	174

This gives a  $4 \times$  improvement in available multiplier count, directly translating to more parallel two-body compute cores on the FPGA fabric.

### 8.2.2 DSP Usage by Arithmetic Submodule

The custom floating-point adder is expected to use approximately zero DSP blocks because it mainly contains exponent comparison, mantissa alignment, add/subtract logic, normalization, and truncation. These operations map primarily to ALMs and registers. The custom floating-point multiplier is

expected to use approximately one DSP block because the main expensive operation is the mantissa multiplication.

The `FastInvSqrt` unit uses a bit-level initial approximation followed by one Newton-Raphson refinement:

$$y \leftarrow y (1.5 - 0.5xy^2).$$

The initial approximation uses approximately zero DSP blocks. The Newton step requires roughly three multiplications:  $y^2$ ,  $xy^2$ , and the final multiplication by  $y$ . Therefore, `FastInvSqrt` is estimated to use approximately three DSP multiplier datapaths. If the following  $f^3$  computation is included,

$$f^2 = f \cdot f, \quad f^3 = f^2 \cdot f,$$

then two additional multipliers are required.

Table 6: Estimated DSP usage of major arithmetic submodules.

Submodule or operation	Estimated DSP multiplier datapaths
Custom floating-point add	0
Custom floating-point multiply	$\approx 1$
Fast inverse square root initial guess	0
Fast inverse square root Newton step	$\approx 3$

A conservative estimate for one fully pipelined `two_body` core is shown in Table 7. This gives approximately 10 DSP multiplier datapaths per core.

Table 7: Estimated DSP multiplier usage per two-body core.

Operation	Estimated multipliers
Squared distance, $dx^2$ and $dy^2$	2
FastInvSqrt Newton step	3
$f^3$ computation	2
Acceleration scaling by $m_j$ , $dx$ , and $dy$	3
Estimated total per two-body core	10

With four parallel `two_body` cores, the expected DSP usage is

$$4 \times 10 = 40$$

DSP multiplier datapaths. This matches the Quartus post-fit report, which shows 40 out of 87 DSP blocks used. The earlier conservative estimate included additional integration multipliers, but the final result suggests that these were either optimized, shared, implemented using logic, or absorbed into existing arithmetic paths.

### 8.2.3 Memory Usage

For the accelerator-side body storage, each body stores five 27-bit values:  $x$ ,  $y$ ,  $v_x$ ,  $v_y$ , and mass. Additional temporary and output storage is needed for intermediate velocity and acceleration values. The conservative accelerator memory estimate is

$$(n \times 5 \times 27) + n + (4 \times n \times 27) + (n \times 4 \times 27) + n = 353n \text{ bits.}$$

For  $n = 1024$ ,

$$353 \times 1024 = 361,472 \text{ bits} \approx 353 \text{ Kbits.}$$

The final Quartus report shows 499,820 block memory bits and 60 RAM blocks used by the complete system. This is larger than the accelerator-only estimate because the fitted design also includes the Avalon-MM peripherals, VGA/display memory, and other system-level storage. However, the total memory utilization is still only 12% of block memory bits and 15% of RAM blocks, so embedded memory is not a limiting resource for this implementation.

## 9 Contributions and Lessons Learned

**Lucy He** - My work focused mainly on the hardware implementation and FPGA integration of the N-body accelerator. I collaborated with Xiyuan and Charlotte on designing and debugging the `two_body_core` and helped extend it into the parallel `four_core_wrapper`. I implemented the leapfrog integrator, participated in the controller FSM design, and helped modify and run testbenches to verify that the hardware outputs matched the Python golden model. I also contributed to the hardware/software memory-map interface, integrated the VGA display interface based on the Lab 3 framework, wrote software tests to check Avalon communication with the accelerator and display peripherals, and handled Quartus compilation, including generating the `.rbf` and device-tree files loaded onto the FPGA. One major lesson I learned came from a difficult bug where fixing the controller made the computed positions correct but unexpectedly broke the VGA display with fractured lines. After checking the code, memory map, software interface, and many possible causes, I concluded that the issue was likely related to post-fit timing or routing changes caused by a small RTL modification, especially because the VGA path is sensitive to stable timing during continuous scanout. This taught me that FPGA debugging goes beyond functional RTL correctness: a design can pass simulation and produce correct numerical outputs while still failing on real hardware due to timing, routing, or integration issues, so post-fit timing analysis and robust system-level testing are essential.

**Xiyuan Peng** - My work mainly focused on the design of the verification infrastructure and part of the hardware implementation for the N-body accelerator. I developed and modified SystemVerilog testbenches for the floating-point arithmetic units, memory modules, controller logic, and accelerator datapath, and helped build Python-based golden-model comparisons to validate the RTL outputs throughout different computation stages. I also participated in the implementation and debugging of several hardware modules, including parts of the two-body compute pipeline, accumulation flow, and controller integration. Through this project, I learned that hardware verification must go beyond module-level functional correctness. Even when simulation outputs match the golden model, FPGA systems can still fail because of timing, routing, interface, or integration issues after synthesis and place-and-route. This project helped me better understand the importance of systematic verification, timing-aware debugging, and full system-level validation in FPGA design.

**Jingzeng Xie** - My work focused mainly on the software implementation and real-time visualization of the N-body accelerator system. I collaborated closely with Pengpeng on most parts of the software stack, including the main application structure, simulation control flow, keyboard interaction, display rendering, and hardware/software communication logic, while Pengpeng primarily handled the accelerator driver. I developed the display-side rendering approach for circular bodies, where precomputed circular stamps are used to efficiently draw particles with different radii onto the framebuffer. This reduced repeated per-pixel computation during each frame and made the software display pipeline more suitable for real-time visualization. I also integrated the software threads for simulation, display, and user input so that the accelerator, VGA output, and keyboard controls could work together as an interactive system. One major lesson I learned is the importance of starting hardware/software integration early. Compared with pure software debugging, FPGA-based projects require much

longer iteration cycles because even small hardware changes may require regeneration, recompilation, device-tree updates, and on-board testing. This made debugging slower and more complex than expected, especially when software behavior depended on hardware timing and interface correctness. Starting integration earlier would leave more time to identify interface mismatches, display issues, and system-level bugs before the final stage of the project.

**Pengpeng Wang** - I implemented the accelerator and display kernel drivers, including the float32-to-27bit conversion and the decision to use a packed 1-bpp framebuffer instead of 32-bpp, which reduced per-frame transfer size from over 1 MB to 38 KB. I also wrote the design document and contributed to the final report. Through this project I gained hands-on experience writing Linux kernel drivers and working with memory-mapped hardware for the first time. Seeing the full stack — from a C program calling `ioctl`, through the kernel, across the Avalon bus, to actual hardware registers — made the abstractions I had only seen in class feel concrete and real.

**Charlotte Chen** My work primarily focuses on the testbench implementation and verification of the accelerator, specifically ensuring the high-precision floating-point, including the FP adders, multipliers, and the fast inverse square root modules, maintained bit-accurate consistency with our Python golden model. For the FP adder and multiplier, I developed directed and constrained-random testbenches to verify corner cases such as subnormal handling, overflow, and rounding modes, ensuring that the pipelined datapath maintained throughput without sacrificing precision. Validating the fast inverse square root was particularly challenging, as it required sweeping a wide range of input magnitudes to ensure the hardware’s approximation stayed within the required error bounds. This taught me that verification must be “precision-aware,” balancing the hardware’s area-efficiency constraints against the numerical stability required for physics simulations. Ultimately, seeing how a minor adjustment in the FISR pipeline could alter the routing and timing of the entire SoC emphasized that in high-performance computing, the physical implementation of arithmetic is just as critical as its logical definition.

## References

- [1] Unknown, “N-body simulation on fpga (iee paper),” *IEEE*, 2018. Available: <https://ieeexplore.ieee.org/document/8445106>.
- [2] id Software, “Quake III arena source code.” <https://github.com/id-Software/Quake-III-Arena>, 1999. Released under GPL v2, 2005.
- [3] F. Stokes, “Fast inverse square root.” <https://github.com/francisrstokes/githubblog/blob/main/2024/5/29/fast-inverse-sqrt.md>, May 2024. Accessed: 2025.
- [4] Wikipedia contributors, “Newton’s method — Wikipedia, the free encyclopedia.” [https://en.wikipedia.org/wiki/Newton%27s\\_method](https://en.wikipedia.org/wiki/Newton%27s_method), 2025. Accessed: 2025.
- [5] Cornell ECE 5760, “Fpga-based n-body simulation final project.” [https://people.ece.cornell.edu/land/courses/ece5760/FinalProjects/s2023/raf269\\_nkg37\\_tjw234/raf269\\_nkg37\\_tjw234/index.html](https://people.ece.cornell.edu/land/courses/ece5760/FinalProjects/s2023/raf269_nkg37_tjw234/raf269_nkg37_tjw234/index.html). Accessed: 2026-04-17.
- [6] Wikipedia, “Leapfrog integration.” [https://en.wikipedia.org/wiki/Leapfrog\\_integration](https://en.wikipedia.org/wiki/Leapfrog_integration). Accessed: 2026-04-17.
- [7] Intel (Altera), “Cyclone v device handbook volume 1.” <https://docs.altera.com/r/docs/683375/current/cyclone-v-device-handbook-volume-1-device-interfaces-and-integration/resources>. Accessed: 2026-04-17.

# A Code

## A.1 Software Source Code

### A.1.1 main.c

Listing 3: code/Software/main.c

```
1 #include <errno.h>
2 #include <pthread.h>
3 #include <stdio.h>
4 #include <stdlib.h>
5 #include <time.h>
6
7 #include "display.h"
8 #include "keyboard.h"
9 #include "nbody.h"
10
11 int main(int argc, char **argv)
12 {
13     pthread_t accelerator_thread;
14     pthread_t vga_thread;
15     pthread_t kb_thread;
16     int have_keyboard = 0;
17
18     // Check command-line arguments
19     if (argc != 3) {
20         fprintf(stderr, "Usage: %s <Num Bodies> <Gap>\n", argv[0]);
21         return 1;
22     }
23
24     // Setting num_bodies and current_gap
25     num_bodies = atoi(argv[1]);
26     current_gap = atoi(argv[2]);
27     if (num_bodies < 5 || num_bodies > MAX_BODIES) {
28         fprintf(stderr, "Num Bodies must be between 5 and %d\n", MAX_BODIES);
29         return 1;
30     }
31     if (current_gap < 1 || current_gap > NBODY_GAP_MAX) {
32         fprintf(stderr, "Gap must be between 1 and %d\n", NBODY_GAP_MAX);
33         return 1;
34     }
35
36     // Allocate history buffer frames with body positions
37     if (allocate_history() < 0) {
38         fprintf(stderr, "Could not allocate history ring (%d frames)\n", MAX_HISTORY);
39         return 1;
40     }
41
42     // Initialize random seed
43     srand((unsigned int) time(NULL));
44
45     // Create Nbody (Accelerator Simulation) Thread
46     if (pthread_create(&accelerator_thread, NULL, nbody_thread, NULL) != 0) {
47         perror("pthread_create nbody_thread");
48         free_history();
49         return 1;
50     }
51
52     // Create Display Thread
53     if (pthread_create(&vga_thread, NULL, display_thread, NULL) != 0) {
54         perror("pthread_create display_thread");
55         pthread_mutex_lock(&state_mutex);
56         running = 0;
57         //pthread_cond_broadcast(&frame_ready_cond);
58         pthread_mutex_unlock(&state_mutex);
59         pthread_join(accelerator_thread, NULL);
60         free_history();
61         return 1;
```

```

62     }
63
64     // Open keyboard device and create Keyboard Thread
65     if (keyboard_open())
66         have_keyboard = 1;
67     if (have_keyboard && pthread_create(&kb_thread, NULL, keyboard_handler, NULL) != 0) {
68         perror("pthread_create keyboard_handler");
69         have_keyboard = 0;
70         free_history();
71         return 1;
72     }
73
74     // Waiting Nbody thread to be destory
75     pthread_join(accelerator_thread, NULL);
76
77     pthread_mutex_lock(&state_mutex);
78     running = 0;
79     //pthread_cond_broadcast(&frame_ready_cond);
80     pthread_mutex_unlock(&state_mutex);
81
82     // Waiting Keyboard and Display thread to be destory
83     pthread_join(kb_thread, NULL);
84     pthread_join(vga_thread, NULL);
85
86     // Clear the history frame buffer
87     free_history();
88
89     return 0;
90 }

```

### A.1.2 nbody.c

Listing 4: code/Software/nbody.c

```

1  #include <errno.h>
2  #include <fcntl.h>
3  #include <pthread.h>
4  #include <stdint.h>
5  #include <stdio.h>
6  #include <stdlib.h>
7  #include <sys/ioctl.h>
8  #include <unistd.h>
9
10 #include "nbody.h"
11
12 int num_bodies = 0;
13 int current_gap = 1;
14 int is_paused = 0;
15 int running = 1;
16
17 uint32_t static_masses[MAX_BODIES];
18 body_pos_t *history[MAX_HISTORY];
19 int h_head = 0;
20 int h_count = 0;
21 int view_idx = 0;
22
23 // Used to distinguish between old and new simulations following a reset
24 static unsigned int sim_generation = 0;
25 // Used to pause after a reset, awaiting user input to continue
26 static int reset_waiting_for_play = 0;
27
28 int nbody_fd = -1;
29
30 pthread_mutex_t state_mutex = PTHREAD_MUTEX_INITIALIZER;
31 pthread_cond_t frame_ready_cond = PTHREAD_COND_INITIALIZER;
32 static pthread_mutex_t hw_mutex = PTHREAD_MUTEX_INITIALIZER;
33
34 // Allocate memory for history frames with body positions
35 int allocate_history(void)

```

```

36 {
37     for (int i = 0; i < MAX_HISTORY; i++) {
38         history[i] = malloc(MAX_BODIES * sizeof(body_pos_t));
39         if (!history[i]) {
40             free_history();
41             return -1;
42         }
43     }
44     return 0;
45 }
46
47 // Release the memory for all history frames
48 void free_history(void)
49 {
50     for (int i = 0; i < MAX_HISTORY; i++) {
51         free(history[i]);
52         history[i] = NULL;
53     }
54 }
55
56 // Generate a random float within the range [min, max]
57 static float random_range(float min, float max)
58 {
59     return min + (max - min) * ((float)rand() / (float)RAND_MAX);
60 }
61
62 // Map the mass to a radius index for display purposes
63 int get_radius_idx(float mass)
64 {
65     float normalized = (mass - NBODY_MASS_MIN) / (NBODY_MASS_MAX - NBODY_MASS_MIN);
66
67     if (normalized < 0.25f)
68         return 0;
69     if (normalized < 0.50f)
70         return 1;
71     if (normalized < 0.75f)
72         return 2;
73     return 3;
74 }
75
76 // Update current_gap in the software, and REG_GAP/REG_N_BODIES in the hardware
77 void nbody_set_gap_delta(int delta)
78 {
79     nbody_config_t cfg;
80
81     // Modify current_gap, then save the new configuration to cfg
82     pthread_mutex_lock(&state_mutex);
83     // If W is pressed, delta > 0, and current_gap is increased
84     if (delta > 0 && current_gap < NBODY_GAP_MAX) {
85         current_gap++;
86     }
87     // If S is pressed, delta < 0, and current_gap is decreased
88     else if (delta < 0 && current_gap > 1) {
89         current_gap--;
90     }
91
92     cfg.num_bodies = (uint32_t) num_bodies;
93     cfg.gap = (uint32_t) current_gap;
94     pthread_mutex_unlock(&state_mutex);
95
96     // Write the new configuration to the hardware device /dev/nbody using ioctl
97     if (nbody_fd >= 0) {
98         pthread_mutex_lock(&hw_mutex);
99         ioctl(nbody_fd, NBODY_WRITE_CONFIG, &cfg);
100        pthread_mutex_unlock(&hw_mutex);
101    }
102 }
103
104 // Used for browsing history frames
105 void nbody_show_frame_delta(int delta)
106 {
107     pthread_mutex_lock(&state_mutex);
108     // User manually advances the frame, the simulation automatically pauses

```

```

109     is_paused = 1;
110
111     // If A is pressed, delta < 0, and the previous frame is displayed
112     if (delta < 0 && view_idx > 0) {
113         view_idx--;
114     }
115     // If D is pressed, delta > 0, and the next frame is displayed
116     else if (delta > 0 && view_idx < h_count - 1) {
117         view_idx++;
118     }
119     //pthread_cond_broadcast(&frame_ready_cond);
120     pthread_mutex_unlock(&state_mutex);
121 }
122
123 /*
124 Function for SPACE (Play/Pause) being pressed
125 Status: Play -> Paused / Paused -> Play
126 */
127 void nbody_toggle_pause(void)
128 {
129     pthread_mutex_lock(&state_mutex);
130     is_paused = !is_paused;
131     // Clear reset waiting status once user switches from the paused to the play
132     if (!is_paused)
133         reset_waiting_for_play = 0;
134     //pthread_cond_broadcast(&frame_ready_cond);
135     pthread_mutex_unlock(&state_mutex);
136 }
137
138 // Function for R (Reset) being pressed
139 void nbody_request_reset(void)
140 {
141     pthread_mutex_lock(&state_mutex);
142     // If a reset has already been performed and the system is waiting for
143     // the user to press Play, do not perform another reset
144     if (reset_waiting_for_play) {
145         pthread_mutex_unlock(&state_mutex);
146         return;
147     }
148     // The simulation automatically pauses
149     is_paused = 1;
150     reset_waiting_for_play = 1;
151     // Revert old results from prior to a reset from being written back to history
152     sim_generation++;
153     //pthread_cond_broadcast(&frame_ready_cond);
154     pthread_mutex_unlock(&state_mutex);
155
156     // Regenerate initial particles and write to hardware
157     reset_system();
158 }
159
160 // Function for Q (Exit) being pressed
161 void nbody_request_quit(void)
162 {
163     pthread_mutex_lock(&state_mutex);
164     // All threads monitor this variable in their main loops
165     // Consequently, this causes the entire program to exit
166     running = 0;
167     //pthread_cond_broadcast(&frame_ready_cond);
168     pthread_mutex_unlock(&state_mutex);
169 }
170
171 // Responsible for initializing a new round of simulation
172 void reset_system(void)
173 {
174     nbody_particle_t *init_particles;
175
176     // Copy global variables to local variables
177     pthread_mutex_lock(&state_mutex);
178     int local_num = num_bodies;
179     int local_gap = current_gap;
180     pthread_mutex_unlock(&state_mutex);
181

```

```

182 // Allocate the initial particle array
183 init_particles = malloc(local_num * sizeof(*init_particles));
184 if (!init_particles) {
185     fprintf(stderr, "failed to allocate initial particles\n");
186     return;
187 }
188
189 // Randomly Generate Initial State
190 pthread_mutex_lock(&state_mutex);
191 for (int i = 0; i < local_num; i++) {
192     // Randomly generate x and y coordinates
193     init_particles[i].x = random_range(NBODY_POS_MIN, NBODY_POS_MAX);
194     init_particles[i].y = random_range(NBODY_POS_MIN, NBODY_POS_MAX);
195     // Save the initial position to history
196     history[0][i].x = init_particles[i].x;
197     history[0][i].y = init_particles[i].y;
198
199     // Randomly generate mass
200     float mass = random_range(NBODY_MASS_MIN, NBODY_MASS_MAX);
201     init_particles[i].mass = mass;
202     // Determine the display radius based on the mass static_masses
203     static_masses[i] = (uint32_t) get_radius_idx(mass);
204
205     // Initialize vx and vy to 0
206     init_particles[i].vx = 0.0f;
207     init_particles[i].vy = 0.0f;
208 }
209 pthread_mutex_unlock(&state_mutex);
210
211 // Write the configuration and initial bodies to the hardware
212 if (nbody_fd >= 0) {
213     nbody_config_t cfg = {
214         .num_bodies = (uint32_t)local_num,
215         .gap = (uint32_t)local_gap,
216     };
217     nbody_bodies_arg_t barg = {
218         .particles = init_particles,
219         .count = (uint32_t)local_num,
220     };
221
222     pthread_mutex_lock(&hw_mutex);
223     // Clear read status / output pointer
224     ioctl(nbody_fd, NBODY_CLEAR_READ);
225     // num_bodies gap
226     if (ioctl(nbody_fd, NBODY_WRITE_CONFIG, &cfg) < 0)
227         perror("NBODY_WRITE_CONFIG");
228     // Write all initial particles to the hardware input buffer
229     if (ioctl(nbody_fd, NBODY_WRITE_BODIES, &barg) < 0)
230         perror("NBODY_WRITE_BODIES");
231     pthread_mutex_unlock(&hw_mutex);
232 }
233
234 // Reset history state
235 pthread_mutex_lock(&state_mutex);
236 // history[0] stores the initial frame and the next written to history[1]
237 h_head = 1;
238 h_count = 1;
239 view_idx = 0;
240 //pthread_cond_broadcast(&frame_ready_cond);
241 pthread_mutex_unlock(&state_mutex);
242
243 // Free the memory for array
244 free(init_particles);
245 }
246
247 // Nbody (Accelerator Simulation) Thread
248 void *nbody_thread(void *arg)
249 {
250     // It is merely there to avoid a compiler warning
251     (void) arg;
252
253     // Open /dev/nbody to obtain a file descriptor
254     // User-space applications communicate with it via ioctl

```

```

255 nbody_fd = open("/dev/nbody", O_RDWR);
256 if (nbody_fd < 0)
257     perror("open /dev/nbody");
258
259 // Regenerate initial particles and write to hardware
260 reset_system();
261
262 /*
263 Check if the process is paused; if not, select the next history slot
264 Initiate hardware computation and wait for completion
265 Read the results from the hardware into next_idx and update the history
266 */
267 while (running) {
268     nbody_read_arg_t rarg;
269
270     int done = 0;
271     unsigned int local_generation;
272
273     // Copy global variables to local variables
274     pthread_mutex_lock(&state_mutex);
275     int paused = is_paused;
276     int local_num = num_bodies;
277     // Subsequently determine whether a reset occurred during this round
278     local_generation = sim_generation;
279     pthread_mutex_unlock(&state_mutex);
280
281     // If paused or the hardware is not open, sleep for 10 ms
282     // Then resume checking in the next iteration
283     if (paused || nbody_fd < 0) {
284         usleep(10000);
285         continue;
286     }
287
288     // After writing to the last slot, it wraps back around to the beginning
289     pthread_mutex_lock(&state_mutex);
290     int next_idx = h_head & (MAX_HISTORY - 1);
291     pthread_mutex_unlock(&state_mutex);
292
293     // Have the kernel driver write REG_GO = 1
294     // Trigger the hardware to begin computation
295     pthread_mutex_lock(&hw_mutex);
296     if (ioctl(nbody_fd, NBODY_START_RUN) < 0) {
297         perror("NBODY_START_RUN");
298         pthread_mutex_unlock(&hw_mutex);
299         usleep(10000);
300         continue;
301     }
302
303     // Polling method waiting for done
304     while (running && !done) {
305         if (ioctl(nbody_fd, NBODY_CHECK_DONE, &done) < 0) {
306             perror("NBODY_CHECK_DONE");
307             break;
308         }
309         if (!done)
310             usleep(1000);
311     }
312
313     // If process being exit
314     if (!running) {
315         pthread_mutex_unlock(&hw_mutex);
316         break;
317     }
318
319     // Once the hardware operations are complete, read the results
320     rarg.results = history[next_idx];
321     rarg.count = (uint32_t) local_num;
322     if (ioctl(nbody_fd, NBODY_READ_RESULTS, &rarg) < 0)
323         perror("NBODY_READ_RESULTS");
324
325     // Clear the read state and prepare for the next read
326     if (ioctl(nbody_fd, NBODY_CLEAR_READ) < 0)
327         perror("NBODY_CLEAR_READ");

```

```

328     pthread_mutex_unlock(&hw_mutex);
329
330     // Update history
331     pthread_mutex_lock(&state_mutex);
332     /*
333     If no reset occurs during this hardware computation cycle,
334     increment normally to the next frame. The old result will not be
335     updated to the history, preventing the old frame position prior
336     to the reset from being marked
337     */
338     if (local_generation == sim_generation) {
339         h_head = (h_head + 1) & (MAX_HISTORY - 1);
340         if (h_count < MAX_HISTORY)
341             h_count++;
342         view_idx = h_count - 1;
343         //pthread_cond_broadcast(&frame_ready_cond);
344     }
345     pthread_mutex_unlock(&state_mutex);
346 }
347
348 // Clear hardware status
349 if (nbody_fd >= 0) {
350     pthread_mutex_lock(&hw_mutex);
351     ioctl(nbody_fd, NBODY_CLEAR_READ);
352     ioctl(nbody_fd, NBODY_STOP);
353     pthread_mutex_unlock(&hw_mutex);
354     close(nbody_fd);
355     nbody_fd = -1;
356 }
357
358 return NULL;
359 }

```

### A.1.3 nbody.h

Listing 5: code/Software/nbody.h

```

1  #ifndef _NBODY_H
2  #define _NBODY_H
3
4  #include <pthread.h>
5  #include <stdint.h>
6  #include "nbody_ioctl.h"
7
8  #define MAX_BODIES NBODY_MAX_BODIES
9  #define MAX_HISTORY 32768
10 #define BODY_DISPLAY_W 640
11 #define BODY_DISPLAY_H 448
12 #define NBODY_POS_MIN (-10.0f)
13 #define NBODY_POS_MAX 10.0f
14 #define NBODY_MASS_MIN 0.002f
15 #define NBODY_MASS_MAX 0.01f
16 #define NBODY_GAP_MAX 10
17
18 extern int num_bodies;
19 extern int current_gap;
20 extern int is_paused;
21 extern int running;
22
23 extern uint32_t static_masses[MAX_BODIES];
24 extern body_pos_t *history[MAX_HISTORY];
25 extern int h_head;
26 extern int h_count;
27 extern int view_idx;
28
29 extern int nbody_fd;
30
31 extern pthread_mutex_t state_mutex;
32 extern pthread_cond_t frame_ready_cond;

```

```

33
34 int allocate_history(void);
35 void free_history(void);
36 int get_radius_idx(float mass);
37 void nbody_set_gap_delta(int delta);
38 void nbody_show_frame_delta(int delta);
39 void nbody_toggle_pause(void);
40 void nbody_request_reset(void);
41 void nbody_request_quit(void);
42 void reset_system(void);
43 void *nbody_thread(void *arg);
44
45 #endif

```

#### A.1.4 display.c

Listing 6: code/Software/display.c

```

1 #include <ctype.h>
2 #include <fcntl.h>
3 #include <stdint.h>
4 #include <stdio.h>
5 #include <stdlib.h>
6 #include <string.h>
7 #include <sys/ioctl.h>
8 #include <unistd.h>
9
10 #include "display.h"
11 #include "nbody.h"
12
13 static int display_fd = -1;
14 static uint32_t framebuffer[DISPLAY_WORDS];
15 static uint8_t body_masks[NUM_BODY_RADII][9][9];
16
17 static const int BODY_RADII[NUM_BODY_RADII] = { 1, 2, 3, 4 };
18
19 // Set a specific pixel is on or off (1/0)
20 static inline void set_pixel(int x, int y, int on)
21 {
22     // Check if the coordinates are out of bounds
23     if (x < 0 || x >= DISPLAY_WIDTH || y < 0 || y >= DISPLAY_HEIGHT)
24         return;
25
26     /*
27     Packed display format
28     frame index = y * 20 + x / 32, bit = x & 31 (x % 32)
29     one LSB-first bit per pixel
30     */
31     if (on) {
32         framebuffer[y * DISPLAY_WORDS_PER_ROW + x / 32] |= (1u << (x & 31));
33     } else {
34         framebuffer[y * DISPLAY_WORDS_PER_ROW + x / 32] &= ~(1u << (x & 31));
35     }
36 }
37
38 // Returning the 5-bit character set of a given character in a given line
39 static uint8_t glyph_row(char c, int row)
40 {
41     // Each character consists of 7 rows, with each row represented by 5 bits
42     static const uint8_t blank[7] = { 0, 0, 0, 0, 0, 0, 0 };
43     static const uint8_t glyphs[][7] = {
44         ['0'] = { 0x0e, 0x11, 0x13, 0x15, 0x19, 0x11, 0x0e },
45         ['1'] = { 0x04, 0x0c, 0x04, 0x04, 0x04, 0x04, 0x0e },
46         ['2'] = { 0x0e, 0x11, 0x01, 0x02, 0x04, 0x08, 0x1f },
47         ['3'] = { 0x1e, 0x01, 0x01, 0x0e, 0x01, 0x01, 0x1e },
48         ['4'] = { 0x02, 0x06, 0x0a, 0x12, 0x1f, 0x02, 0x02 },
49         ['5'] = { 0x1f, 0x10, 0x10, 0x1e, 0x01, 0x01, 0x1e },
50         ['6'] = { 0x0e, 0x10, 0x10, 0x1e, 0x11, 0x11, 0x0e },
51         ['7'] = { 0x1f, 0x01, 0x02, 0x04, 0x08, 0x08, 0x08 },

```

```

52     ['8'] = { 0x0e, 0x11, 0x11, 0x0e, 0x11, 0x11, 0x0e },
53     ['9'] = { 0x0e, 0x11, 0x11, 0x0f, 0x01, 0x01, 0x0e },
54     ['A'] = { 0x0e, 0x11, 0x11, 0x1f, 0x11, 0x11, 0x11 },
55     ['B'] = { 0x1e, 0x11, 0x11, 0x1e, 0x11, 0x11, 0x1e },
56     ['C'] = { 0x0f, 0x10, 0x10, 0x10, 0x10, 0x10, 0x0f },
57     ['D'] = { 0x1e, 0x11, 0x11, 0x11, 0x11, 0x11, 0x1e },
58     ['E'] = { 0x1f, 0x10, 0x10, 0x1e, 0x10, 0x10, 0x1f },
59     ['F'] = { 0x1f, 0x10, 0x10, 0x1e, 0x10, 0x10, 0x10 },
60     ['G'] = { 0x0f, 0x10, 0x10, 0x13, 0x11, 0x11, 0x0f },
61     ['H'] = { 0x11, 0x11, 0x11, 0x1f, 0x11, 0x11, 0x11 },
62     ['I'] = { 0x0e, 0x04, 0x04, 0x04, 0x04, 0x04, 0x0e },
63     ['J'] = { 0x07, 0x02, 0x02, 0x02, 0x12, 0x12, 0x0c },
64     ['K'] = { 0x11, 0x12, 0x14, 0x18, 0x14, 0x12, 0x11 },
65     ['L'] = { 0x10, 0x10, 0x10, 0x10, 0x10, 0x10, 0x1f },
66     ['M'] = { 0x11, 0x1b, 0x15, 0x15, 0x11, 0x11, 0x11 },
67     ['N'] = { 0x11, 0x19, 0x15, 0x13, 0x11, 0x11, 0x11 },
68     ['O'] = { 0x0e, 0x11, 0x11, 0x11, 0x11, 0x11, 0x0e },
69     ['P'] = { 0x1e, 0x11, 0x11, 0x1e, 0x10, 0x10, 0x10 },
70     ['Q'] = { 0x0e, 0x11, 0x11, 0x11, 0x15, 0x12, 0x0d },
71     ['R'] = { 0x1e, 0x11, 0x11, 0x1e, 0x14, 0x12, 0x11 },
72     ['S'] = { 0x0f, 0x10, 0x10, 0x0e, 0x01, 0x01, 0x1e },
73     ['T'] = { 0x1f, 0x04, 0x04, 0x04, 0x04, 0x04, 0x04 },
74     ['U'] = { 0x11, 0x11, 0x11, 0x11, 0x11, 0x11, 0x0e },
75     ['V'] = { 0x11, 0x11, 0x11, 0x11, 0x11, 0x0a, 0x04 },
76     ['W'] = { 0x11, 0x11, 0x11, 0x15, 0x15, 0x15, 0x0a },
77     ['X'] = { 0x11, 0x11, 0x0a, 0x04, 0x0a, 0x11, 0x11 },
78     ['Y'] = { 0x11, 0x11, 0x0a, 0x04, 0x04, 0x04, 0x04 },
79     ['Z'] = { 0x1f, 0x01, 0x02, 0x04, 0x08, 0x10, 0x1f },
80     [':'] = { 0x00, 0x04, 0x04, 0x00, 0x04, 0x04, 0x00 },
81     ['/'] = { 0x01, 0x01, 0x02, 0x04, 0x08, 0x10, 0x10 },
82     ['['] = { 0x0e, 0x08, 0x08, 0x08, 0x08, 0x08, 0x0e },
83     [']'] = { 0x0e, 0x02, 0x02, 0x02, 0x02, 0x02, 0x0e },
84     ['|'] = { 0x04, 0x04, 0x04, 0x04, 0x04, 0x04, 0x04 },
85     ['-'] = { 0x00, 0x00, 0x00, 0x1f, 0x00, 0x00, 0x00 },
86     [' '] = { 0x00, 0x00, 0x00, 0x00, 0x00, 0x00, 0x00 },
87 };
88 const uint8_t *g;
89
90 // Converts lowercase letters to uppercase
91 if (c >= 'a' && c <= 'z')
92     c = (char)toupper((unsigned char)c);
93
94 // If the character is not supported, display a blank space
95 if ((unsigned char) c >= sizeof(glyphs) / sizeof(glyphs[0])) {
96     g = blank;
97 } else {
98     g = glyphs[(unsigned char) c];
99 }
100
101 return g[row];
102 }
103
104 // Convert the simulation's world coordinates into screen pixel coordinates
105 int world_to_screen_x(float x)
106 {
107     float normalized = (x - NBODY_POS_MIN) / (NBODY_POS_MAX - NBODY_POS_MIN);
108     return (int) (normalized * (float)(BODY_DISPLAY_W - 1) + 0.5f);
109 }
110
111 // Convert the simulation's world coordinates into screen pixel coordinates
112 int world_to_screen_y(float y)
113 {
114     float normalized = (y - NBODY_POS_MIN) / (NBODY_POS_MAX - NBODY_POS_MIN);
115     return (int) (normalized * (float)(BODY_DISPLAY_H - 1) + 0.5f);
116 }
117
118 // Clear the entire framebuffer to 0, resulting in a black screen
119 void display_clear(void)
120 {
121     memset(framebuffer, 0, sizeof(framebuffer));
122 }
123
124 /*

```

```

125 The rendering thread needs to draw numerous bodies in every frame
126 Recalculating the circular geometry each time a body is drawn is time-consuming
127 Instead, pre-generate the masks beforehand
128 then simply perform a direct lookup when rendering
129 */
130 void display_init_bodyshape(void)
131 {
132     memset(body_masks, 0, sizeof(body_masks));
133
134     for (int i = 0; i < NUM_BODY_RADII; i++) {
135         int r = BODY_RADII[i];
136         for (int y = -r; y <= r; y++)
137             for (int x = -r; x <= r; x++)
138                 if (x * x + y * y < r * r + r)
139                     body_masks[i][y + r][x + r] = 1;
140     }
141 }
142
143 // Draw a body at position (cx, cy)
144 void display_draw_body(int cx, int cy, int radius_idx)
145 {
146     int r = BODY_RADII[radius_idx];
147     for (int y = -r; y <= r; y++) {
148         int sy = cy + y;
149
150         // Check if the coordinates are out of bounds
151         if (sy < 0 || sy >= BODY_DISPLAY_H)
152             continue;
153         for (int x = -r; x <= r; x++)
154             if (body_masks[radius_idx][y + r][x + r])
155                 set_pixel(cx + x, sy, 1);
156     }
157 }
158
159 // Draw a character at a specific row and column within a text grid
160 void display_putchar(char c, int row, int col)
161 {
162     // Check if the coordinates are out of bounds
163     if (row < 0 || row >= TEXT_ROWS || col < 0 || col >= TEXT_COLS)
164         return;
165
166     // Render 5 x 7 character to 8 16 pixels
167     for (int y = 0; y < FONT_HEIGHT; y++) {
168         uint8_t bits = 0;
169
170         if (y >= 2 && y < 16)
171             bits = glyph_row(c, (y - 2) / 2);
172
173         for (int x = 0; x < FONT_WIDTH; x++) {
174             int on = (x >= 1 && x <= 5) && (bits & (1u << (5 - x)));
175             set_pixel(col * FONT_WIDTH + x, row * FONT_HEIGHT + y, on);
176         }
177     }
178 }
179
180 // Draw Strings Continuously
181 void display_puts(const char *s, int row, int col)
182 {
183     while (*s && col < TEXT_COLS) {
184         display_putchar(*s++, row, col++);
185     }
186 }
187
188 // Send the software framebuffer to the kernel driver
189 // If /dev/nbody_display is not open, return -1
190 int display_present(void)
191 {
192     if (display_fd < 0)
193         return -1;
194     return ioctl(display_fd, DISPLAY_WRITE_FRAME, framebuffer);
195 }
196
197 // Display Thread

```

```

198 void *display_thread(void *arg)
199 {
200     // It is merely there to avoid a compiler warning
201     (void) arg;
202
203     // Allocate a temporary array to store the positions of the current frame
204     body_pos_t *render_positions = malloc(MAX_BODIES * sizeof(*render_positions));
205     if (!render_positions)
206         return NULL;
207
208     // Open display device
209     display_fd = open("/dev/nbody_display", O_RDWR);
210     if (display_fd < 0)
211         perror("open /dev/nbody_display");
212
213     // Initialize circular mask for the body
214     display_init_bodyshape();
215
216     while (running) {
217         // Copy a local variables from the global variables
218         pthread_mutex_lock(&state_mutex);
219         int local_num = num_bodies;
220         int local_gap = current_gap;
221         int local_count = h_count;
222         int local_view = view_idx;
223         int local_paused = is_paused;
224
225         // Select the current frame to be displayed from history frames
226         if (local_count > 0) {
227             int slot = (local_count < MAX_HISTORY) ? local_view : (h_head + local_view) & (
MAX_HISTORY - 1);
228             memcpy(render_positions, history[slot], local_num * sizeof(*render_positions));
229         }
230         pthread_mutex_unlock(&state_mutex);
231
232         // Clear the display framebuffer
233         display_clear();
234
235         // If historical frames exist
236         if (local_count > 0) {
237             for (int i = 0; i < local_num; i++) {
238                 // Retrieve world x/y coordinates and convert to screen x/y coordinates
239                 int x = world_to_screen_x(render_positions[i].x);
240                 int y = world_to_screen_y(render_positions[i].y);
241                 display_draw_body(x, y, (int)static_masses[i]);
242             }
243         }
244
245         // Display the UI in last two rows in screen/VGA
246         char line1[TEXT_COLS + 1];
247         snprintf(line1, sizeof(line1), "Bodies: %4d/%d | Gap: %2d/%d | Frame: %d/%d | Status
: %s",
248             local_num, MAX_BODIES, local_gap, NBODY_GAP_MAX, local_view + 1,
local_count,
249             local_paused ? "PAUSED" : "RUNNING");
250         display_puts(line1, UI_START_ROW, 0);
251         display_puts("[SPACE] Play/Pause | [W/S] Gap | [A/D] Frame | [R] Reset | [Q] Exit",
252             UI_START_ROW + 1, 0);
253
254         // Write the entire framebuffer to the display hardware
255         display_present();
256
257         // It sleeps for 33.333 ms, the display thread is running at approximately 30 fps
258         usleep(33333);
259     }
260
261     // Clear the screen once before exiting and close the display device
262     if (display_fd >= 0) {
263         display_clear();
264         display_present();
265         close(display_fd);
266         display_fd = -1;
267     }

```

```

268
269     // Release temporary array
270     free(render_positions);
271
272     return NULL;
273 }

```

### A.1.5 display.h

Listing 7: code/Software/display.h

```

1  #ifndef _DISPLAY_H
2  #define _DISPLAY_H
3
4  #include "display_ioctl.h"
5
6  #define FONT_WIDTH 8
7  #define FONT_HEIGHT 16
8  #define TEXT_COLS 80
9  #define TEXT_ROWS 30
10 #define UI_START_ROW 28
11 #define BODY_DISPLAY_H 448
12 #define NUM_BODY_RADII 4
13
14 int world_to_screen_x(float x);
15 int world_to_screen_y(float y);
16 void display_clear(void);
17 void display_init_bodyshape(void);
18 void display_draw_body(int cx, int cy, int radius_idx);
19 void display_putchar(char c, int row, int col);
20 void display_puts(const char *s, int row, int col);
21 int display_present(void);
22 void *display_thread(void *arg);
23
24 #endif

```

### A.1.6 keyboard.c

Listing 8: code/Software/keyboard.c

```

1  #include <stdint.h>
2  #include <stdio.h>
3  #include <string.h>
4  #include <time.h>
5
6  #include "keyboard.h"
7  #include "nbody.h"
8  #include "usbkeyboard.h"
9
10 static struct libusb_device_handle *keyboard = NULL;
11 static uint8_t endpoint_address;
12
13 // Open keyboard device
14 int keyboard_open(void)
15 {
16     // This function is implemented in usbkeyboard.c
17     // it scans USB devices to locate a HID keyboard
18     keyboard = openkeyboard(&endpoint_address);
19     if (!keyboard)
20         fprintf(stderr, "Warning: did not find a USB keyboard\n");
21
22     return keyboard != NULL;
23 }
24
25 /*

```

```

26 This function returns the current monotonic time in milliseconds
27 It uses CLOCK_MONOTONIC rather than standard wall-clock time
28 The monotonic clock does not jump or shift when the system time is adjusted
29 More suitable for timing key-repeat intervals
30 */
31 static unsigned long long monotonic_ms(void)
32 {
33     struct timespec ts;
34
35     clock_gettime(CLOCK_MONOTONIC, &ts);
36     return (unsigned long long)ts.tv_sec * 1000ull +
37         (unsigned long long)ts.tv_nsec / 1000000ull;
38 }
39
40 // Converting USB HID Keycodes into Program Actions
41 static void handle_keycode(uint8_t keycode)
42 {
43     // SPACE - Pause/Play
44     if (keycode == 0x2c) {
45         nbody_toggle_pause();
46     }
47     // W - Gap Increase
48     else if (keycode == 0x1a) {
49         nbody_set_gap_delta(1);
50     }
51     // S - Gap Decrease
52     else if (keycode == 0x16) {
53         nbody_set_gap_delta(-1);
54     }
55     // A - Frame Backward
56     else if (keycode == 0x04) {
57         nbody_show_frame_delta(-1);
58     }
59     // D - Frame Forward
60     else if (keycode == 0x07) {
61         nbody_show_frame_delta(1);
62     }
63     // R - Reset
64     else if (keycode == 0x15) {
65         nbody_request_reset();
66     }
67     // Q - Exit
68     else if (keycode == 0x14) {
69         nbody_request_quit();
70     }
71 }
72
73 /* Keyboard Thread
74 Read a packet from the USB keyboard, waiting for a maximum of 50 ms
75 If a timeout occurs: if the 'A' or 'D' key is currently being held down
76 and the repeat interval has elapsed, trigger the corresponding action
77 Otherwise, proceed to the next iteration
78 If the read operation fails, proceed to the next iteration
79 Check if the current packet is identical to the previous packet
80 If no key is currently pressed, clear the repeat state and proceed to the next iteration
81 If the key is A/D, enable long-press repeat functionality and proceed to the next iteration
82 For all other keys: if the current packet is identical to the previous packet, ignore it
83 Otherwise, process the key press once
84 */
85 void *keyboard_handler(void *arg)
86 {
87     // It is merely there to avoid a compiler warning
88     (void) arg;
89
90     /*
91     These two values are used to scroll through historical frames when holding down A/
92     D
93     After the initial press of A/D, a 350 ms delay before continuous repetition begins
94     // Once repetition starts, a frame-scrolling event is triggered every 35 ms
95     */
96     enum {
97         REPEAT_DELAY_MS = 350,
98         REPEAT_INTERVAL_MS = 35,

```

```

98     };
99
100    // Data packets read from a USB keyboard
101    struct usb_keyboard_packet packet;
102    uint8_t last_packet[8] = { 0 };
103    uint8_t held_repeat_key = 0;
104    unsigned long long next_repeat_ms = 0;
105
106    while (running) {
107        /*
108        Read a keyboard packet from the USB keyboard's interrupt endpoint
109        If no keyboard event occurs within 50 ms, it will time out
110        rather than blocking indefinitely
111        */
112        int transferred = 0;
113        int rc = libusb_interrupt_transfer(keyboard, endpoint_address,
114                                         (unsigned char *)&packet, sizeof(packet),
115                                         &transferred, 50);
116        unsigned long long now = monotonic_ms();
117
118        /*
119        If no new keyboard packet is received this time, but a k e y specifically A/D
120        was previously being held down (repeating), then as soon as the designated
121        time interval elapses, trigger the action once again
122        */
123        if (rc == LIBUSB_ERROR_TIMEOUT) {
124            if (held_repeat_key && now >= next_repeat_ms) {
125                handle_keycode(held_repeat_key);
126                next_repeat_ms = now + REPEAT_INTERVAL_MS;
127            }
128            continue;
129        }
130
131        //If a USB transfer error occurs, or if the number of bytes read is incorrect
132        if (rc != 0 || transferred != sizeof(packet))
133            continue;
134
135        // Determines whether the current packet is identical to the previous one
136        // in order to prevent a single key press from being processed repeatedly
137        int same_packet = (memcmp(&packet, last_packet, sizeof(packet)) == 0);
138        if (!same_packet)
139            memcpy(last_packet, &packet, sizeof(packet));
140
141        // Take only the first keycode
142        uint8_t keycode = packet.keycode[0];
143        if (keycode == 0) {
144            held_repeat_key = 0;
145            continue;
146        }
147
148        // A and D are special because they support long-press repeat
149        if (keycode == 0x04 || keycode == 0x07) {
150            if (!same_packet || held_repeat_key != keycode) {
151                handle_keycode(keycode);
152                held_repeat_key = keycode;
153                next_repeat_ms = now + REPEAT_DELAY_MS;
154            } else if (now >= next_repeat_ms) {
155                handle_keycode(keycode);
156                next_repeat_ms = now + REPEAT_INTERVAL_MS;
157            }
158            continue;
159        }
160
161        // If the current packet is identical to the previous one
162        // it indicates that the key is still being held down
163        // then do not trigger a duplicate event
164        held_repeat_key = 0;
165        if (same_packet)
166            continue;
167
168        handle_keycode(keycode);
169    }
170

```

```

171     return NULL;
172 }

```

### A.1.7 keyboard.h

Listing 9: code/Software/keyboard.h

```

1  #ifndef _KEYBOARD_H
2  #define _KEYBOARD_H
3
4  int keyboard_open(void);
5  void *keyboard_handler(void *arg);
6
7  #endif

```

### A.1.8 usbkeyboard.c

Listing 10: code/Software/usbkeyboard.c

```

1  #include "usbkeyboard.h"
2
3  #include <stdio.h>
4  #include <stdlib.h>
5
6  /* References on libusb 1.0 and the USB HID/keyboard protocol
7  *
8  * http://libusb.org
9  * https://web.archive.org/web/20210302095553/https://www.dreamincode.net/forums/topic
   /148707-introduction-to-using-libusb-10/
10 *
11 * https://www.usb.org/sites/default/files/documents/hid1_11.pdf
12 *
13 * https://usb.org/sites/default/files/hut1_5.pdf
14 */
15
16 /*
17 * Find and return a USB keyboard device or NULL if not found
18 * The argument con
19 *
20 */
21 struct libusb_device_handle *openkeyboard(uint8_t *endpoint_address) {
22     libusb_device **devs;
23     struct libusb_device_handle *keyboard = NULL;
24     struct libusb_device_descriptor desc;
25     ssize_t num_devs, d;
26     uint8_t i, k;
27
28     /* Start the library */
29     if ( libusb_init(NULL) < 0 ) {
30         fprintf(stderr, "Error: libusb_init failed\n");
31         exit(1);
32     }
33
34     /* Enumerate all the attached USB devices */
35     if ( (num_devs = libusb_get_device_list(NULL, &devs)) < 0 ) {
36         fprintf(stderr, "Error: libusb_get_device_list failed\n");
37         exit(1);
38     }
39
40     /* Look at each device, remembering the first HID device that speaks
41     the keyboard protocol */
42
43     for (d = 0 ; d < num_devs ; d++) {
44         libusb_device *dev = devs[d];
45         if ( libusb_get_device_descriptor(dev, &desc) < 0 ) {

```

```

46     fprintf(stderr, "Error: libusb_get_device_descriptor failed\n");
47     exit(1);
48 }
49
50 if (desc.bDeviceClass == LIBUSB_CLASS_PER_INTERFACE) {
51     struct libusb_config_descriptor *config;
52     libusb_get_config_descriptor(dev, 0, &config);
53     for (i = 0 ; i < config->bNumInterfaces ; i++)
54     for ( k = 0 ; k < config->interface[i].num_altsetting ; k++ ) {
55         const struct libusb_interface_descriptor *inter =
56             config->interface[i].altsetting + k ;
57         if ( inter->bInterfaceClass == LIBUSB_CLASS_HID &&
58             inter->bInterfaceProtocol == USB_HID_KEYBOARD_PROTOCOL) {
59             int r;
60             if ((r = libusb_open(dev, &keyboard)) != 0) {
61                 fprintf(stderr, "Error: libusb_open failed: %d\n", r);
62                 exit(1);
63             }
64             if (libusb_kernel_driver_active(keyboard,i))
65                 libusb_detach_kernel_driver(keyboard, i);
66             libusb_set_auto_detach_kernel_driver(keyboard, i);
67             if ((r = libusb_claim_interface(keyboard, i)) != 0) {
68                 fprintf(stderr, "Error: libusb_claim_interface failed: %d\n", r);
69                 exit(1);
70             }
71             *endpoint_address = inter->endpoint[0].bEndpointAddress;
72             goto found;
73         }
74     }
75 }
76 }
77
78 found:
79     libusb_free_device_list(devs, 1);
80
81     return keyboard;
82 }

```

### A.1.9 usbkeyboard.h

Listing 11: code/Software/usbkeyboard.h

```

1  #ifndef _USBKEYBOARD_H
2  #define _USBKEYBOARD_H
3
4  #include <libusb-1.0/libusb.h>
5
6  #define USB_HID_KEYBOARD_PROTOCOL 1
7
8  /* Modifier bits */
9  #define USB_LCTRL (1 << 0)
10 #define USB_LSHIFT (1 << 1)
11 #define USB_LALT (1 << 2)
12 #define USB_LGUI (1 << 3)
13 #define USB_RCTRL (1 << 4)
14 #define USB_RSHIFT (1 << 5)
15 #define USB_RALT (1 << 6)
16 #define USB_RGUI (1 << 7)
17
18 struct usb_keyboard_packet {
19     uint8_t modifiers;
20     uint8_t reserved;
21     uint8_t keycode[6];
22 };
23
24 /* Find and open a USB keyboard device. Argument should point to
25 space to store an endpoint address. Returns NULL if no keyboard
26 device was found. */
27 extern struct libusb_device_handle *openkeyboard(uint8_t *);

```

```
28
29 #endif
```

## A.1.10 accelerator\_driver.c

Listing 12: code/Software/accelerator\_driver.c

```
1  /*
2  * N-body Avalon-MM Linux driver.
3  *
4  * This follows nbody_accel_avmm.sv exactly. Register addresses below are
5  * byte offsets for ioread32/iowrite32; the hardware Avalon address is a
6  * 32-bit word address, so every register is word_address * 4.
7  */
8
9  #include <linux/errno.h>
10 #include <linux/fs.h>
11 #include <linux/init.h>
12 #include <linux/io.h>
13 #include <linux/kernel.h>
14 #include <linux/miscdevice.h>
15 #include <linux/module.h>
16 #include <linux/of.h>
17 #include <linux/of_address.h>
18 #include <linux/platform_device.h>
19 #include <linux/slab.h>
20 #include <linux/string.h>
21 #include <linux/uaccess.h>
22 #include "nbody_ioctl.h"
23
24 #define DRIVER_NAME "nbody"
25
26 #define REG_GO          (0x00 * 4)
27 #define REG_N_BODIES   (0x01 * 4)
28 #define REG_GAP        (0x02 * 4)
29 #define REG_X_IN       (0x03 * 4)
30 #define REG_Y_IN       (0x04 * 4)
31 #define REG_M_IN       (0x05 * 4)
32 #define REG_VX_IN      (0x06 * 4)
33 #define REG_VY_IN      (0x07 * 4)
34 #define REG_DONE       (0x08 * 4)
35 #define REG_READ       (0x09 * 4)
36 #define REG_OUT_X      (0x0A * 4)
37 #define REG_OUT_Y      (0x0B * 4)
38
39 struct nbody_dev {
40     struct resource res;
41     void __iomem *virtbase;
42     uint32_t num_bodies;
43     uint32_t gap;
44 };
45
46 static struct nbody_dev dev;
47
48 /*
49 * Convert a userspace IEEE-754 single-precision float into the accelerator's
50 * custom S1E8M18 27-bit floating-point format.
51 *
52 * The converted value is stored in the lower 27 bits of a 32-bit Avalon word.
53 * The upper 5 bits are unused padding.
54 */
55 static uint32_t f32_to_f27_bits(const float *value)
56 {
57     uint32_t bits;
58     uint32_t sign;
59     uint32_t exp;
60     uint32_t mant;
61     uint32_t rounded_mant;
62
```

```

63     memcpy(&bits, value, sizeof(bits));
64
65     sign = (bits >> 31) & 1u;
66     exp = (bits >> 23) & 0xffu;
67     mant = bits & 0x007ffffu;
68
69     if (exp == 0)
70         return 0;
71     if (exp == 0xffu)
72         return (sign << NBODY_F27_SIGN_SHIFT) | (0xfeu << NBODY_F27_EXP_SHIFT) |
73             NBODY_F27_MANT_MASK;
74
75     rounded_mant = (mant + 0x10u) >> 5;
76     if (rounded_mant == (1u << 18)) {
77         rounded_mant = 0;
78         exp++;
79         if (exp >= 0xffu) {
80             exp = 0xfeu;
81             rounded_mant = NBODY_F27_MANT_MASK;
82         }
83     }
84
85     return (sign << NBODY_F27_SIGN_SHIFT) |
86         (exp << NBODY_F27_EXP_SHIFT) |
87         (rounded_mant & NBODY_F27_MANT_MASK);
88 }
89
90 /*
91  * Convert a raw 27-bit S1E8M18 hardware result back into an IEEE-754
92  * single-precision float for userspace.
93  */
94 static void f27_bits_to_f32(uint32_t raw, float *value)
95 {
96     uint32_t bits;
97     uint32_t sign;
98     uint32_t exp;
99     uint32_t mant;
100
101     raw &= NBODY_DATA_MASK;
102     if (raw == 0) {
103         bits = 0;
104     } else {
105         sign = (raw >> NBODY_F27_SIGN_SHIFT) & 1u;
106         exp = (raw >> NBODY_F27_EXP_SHIFT) & 0xffu;
107         mant = raw & NBODY_F27_MANT_MASK;
108
109         if (exp == 0) {
110             bits = 0;
111         } else {
112             bits = (sign << 31) | (exp << 23) | (mant << 5);
113         }
114     }
115
116     memcpy(value, &bits, sizeof(bits));
117 }
118
119 /*
120  * Copy the initial particle array from userspace, convert each field from
121  * float32 to S1E8M18, and stream the data into the hardware input registers.
122  *
123  * The hardware commits one complete body when VY_IN is written, so the fields
124  * must be written in the expected X, Y, mass, VX, VY order.
125  */
126 static int nbody_write_bodies(unsigned long arg)
127 {
128     nbody_bodies_arg_t barg;
129     nbody_particle_t *particles;
130     size_t bytes;
131     uint32_t i;
132
133     if (copy_from_user(&barg, (void __user *)arg, sizeof(barg)))
134         return -EFAULT;
135

```

```

136     if (barg.count == 0 || barg.count > NBODY_MAX_BODIES || !barg.particles)
137         return -EINVAL;
138
139     if (dev.num_bodies == 0 || barg.count != dev.num_bodies)
140         return -EINVAL;
141
142     bytes = barg.count * sizeof(*particles);
143     particles = memdup_user((const void __user *)barg.particles, bytes);
144     if (IS_ERR(particles))
145         return PTR_ERR(particles);
146
147     /*
148      * Userspace passes IEEE-754 single-precision values. Hardware payloads are
149      * rounded to S1E8M18 in the low 27 bits of each 32-bit Avalon word; the
150      * upper 5 bits are zero padding. Writing VY_IN commits the current body
151      * and advances input_ptr.
152      */
153     for (i = 0; i < barg.count; i++) {
154         iowrite32(f32_to_f27_bits(&particles[i].x), dev.virtbase + REG_X_IN);
155         iowrite32(f32_to_f27_bits(&particles[i].y), dev.virtbase + REG_Y_IN);
156         iowrite32(f32_to_f27_bits(&particles[i].mass), dev.virtbase + REG_M_IN);
157         iowrite32(f32_to_f27_bits(&particles[i].vx), dev.virtbase + REG_VX_IN);
158         iowrite32(f32_to_f27_bits(&particles[i].vy), dev.virtbase + REG_VY_IN);
159     }
160
161     kfree(particles);
162     return 0;
163 }
164
165 /*
166  * Read computed body positions from the accelerator output registers.
167  *
168  * Writing READ=1 prepares the hardware output stream. For each body, OUT_X is
169  * read first and OUT_Y is read second. The hardware output pointer advances
170  * only after OUT_Y is read.
171  */
172 static int nbody_read_results(unsigned long arg)
173 {
174     nbody_read_arg_t rarg;
175     nbody_result_t *results;
176     size_t bytes;
177     uint32_t i;
178     int ret = 0;
179
180     if (copy_from_user(&rarg, (void __user *)arg, sizeof(rarg)))
181         return -EFAULT;
182
183     if (rarg.count == 0 || rarg.count > NBODY_MAX_BODIES ||
184         rarg.count > dev.num_bodies || !rarg.results)
185         return -EINVAL;
186
187     bytes = rarg.count * sizeof(*results);
188     results = kmalloc(bytes, GFP_KERNEL);
189     if (!results)
190         return -ENOMEM;
191
192     iowrite32(1, dev.virtbase + REG_READ);
193
194     /*
195      * Read OUT_X first, then OUT_Y. Only the OUT_Y read increments the
196      * hardware output pointer.
197      */
198     for (i = 0; i < rarg.count; i++) {
199         f27_bits_to_f32(ioread32(dev.virtbase + REG_OUT_X), &results[i].x);
200         f27_bits_to_f32(ioread32(dev.virtbase + REG_OUT_Y), &results[i].y);
201     }
202
203     if (copy_to_user((void __user *)rarg.results, results, bytes))
204         ret = -EFAULT;
205
206     kfree(results);
207     return ret;
208 }

```

```

209
210 /*
211  * Main ioctl dispatch function for /dev/nbody.
212  *
213  * Each command corresponds to one userspace accelerator operation, such as
214  * writing configuration, loading input bodies, starting a run, checking DONE,
215  * reading output positions, or clearing the hardware read state.
216  */
217 static long nbody_ioctl(struct file *f, unsigned int cmd, unsigned long arg)
218 {
219     nbody_config_t cfg;
220     int done;
221
222     switch (cmd) {
223     case NBODY_WRITE_CONFIG:
224         if (copy_from_user(&cfg, (void __user *)arg, sizeof(cfg)))
225             return -EFAULT;
226         if (cfg.num_bodies > NBODY_MAX_BODIES || cfg.gap == 0)
227             return -EINVAL;
228
229         dev.num_bodies = cfg.num_bodies;
230         dev.gap = cfg.gap;
231         iowrite32(cfg.num_bodies, dev.virtbase + REG_N_BODIES);
232         iowrite32(cfg.gap, dev.virtbase + REG_GAP);
233         return 0;
234
235     case NBODY_WRITE_BODIES:
236         return nbody_write_bodies(arg);
237
238     case NBODY_START_RUN:
239         iowrite32(1, dev.virtbase + REG_GO);
240         return 0;
241
242     case NBODY_CHECK_DONE:
243         done = ioread32(dev.virtbase + REG_DONE) & 1;
244         if (copy_to_user((int __user *)arg, &done, sizeof(done)))
245             return -EFAULT;
246         return 0;
247
248     case NBODY_READ_RESULTS:
249         return nbody_read_results(arg);
250
251     case NBODY_CLEAR_READ:
252         iowrite32(0, dev.virtbase + REG_READ);
253         return 0;
254
255     case NBODY_STOP:
256         /* GO is pulse-based in SV; best-effort stop is clearing READ. */
257         iowrite32(0, dev.virtbase + REG_READ);
258         return 0;
259
260     case NBODY_SOFT_RESET:
261         /*
262          * There is no full software reset register. Clear READ/output_ptr;
263          * userspace must reload config and bodies for a fresh simulation.
264          */
265         iowrite32(0, dev.virtbase + REG_READ);
266         return 0;
267
268     default:
269         return -EINVAL;
270     }
271 }
272
273 static const struct file_operations nbody_fops = {
274     .owner = THIS_MODULE,
275     .unlocked_ioctl = nbody_ioctl,
276 };
277
278 static struct miscdevice nbody_misc_device = {
279     .minor = MISC_DYNAMIC_MINOR,
280     .name = "nbody",
281     .fops = &nbody_fops,

```

```

282 };
283
284 /*
285  * Platform driver probe function.
286  *
287  * This function is called when the matching device-tree node is found. It maps
288  * the accelerator register space, registers /dev/nbody as a misc device, and
289  * initializes basic hardware control registers.
290  */
291 static int nbody_probe(struct platform_device *pdev)
292 {
293     int ret;
294
295     ret = of_address_to_resource(pdev->dev.of_node, 0, &dev.res);
296     if (ret)
297         return ret;
298
299     if (!request_mem_region(dev.res.start, resource_size(&dev.res), DRIVER_NAME))
300         return -EBUSY;
301
302     dev.virtbase = of_iomap(pdev->dev.of_node, 0);
303     if (!dev.virtbase) {
304         ret = -ENOMEM;
305         goto out_release;
306     }
307
308     ret = misc_register(&nbody_misc_device);
309     if (ret)
310         goto out_unmap;
311
312     iowrite32(0, dev.virtbase + REG_READ);
313     iowrite32(1, dev.virtbase + REG_GAP);
314     pr_info("nbody driver registered /dev/nbody\n");
315     return 0;
316
317 out_unmap:
318     iounmap(dev.virtbase);
319 out_release:
320     release_mem_region(dev.res.start, resource_size(&dev.res));
321     return ret;
322 }
323
324 /*
325  * Platform driver remove function.
326  *
327  * Release the misc device, unmap the hardware register space, and free the
328  * reserved memory region.
329  */
330 static int nbody_remove(struct platform_device *pdev)
331 {
332     misc_deregister(&nbody_misc_device);
333     if (dev.virtbase)
334         iounmap(dev.virtbase);
335     release_mem_region(dev.res.start, resource_size(&dev.res));
336     return 0;
337 }
338
339 static const struct of_device_id nbody_of_match[] = {
340     { .compatible = "csee4840,nbody_accel_avmm-1.0" },
341     {}
342 };
343 MODULE_DEVICE_TABLE(of, nbody_of_match);
344
345 static struct platform_driver nbody_driver = {
346     .probe = nbody_probe,
347     .remove = nbody_remove,
348     .driver = {
349         .name = DRIVER_NAME,
350         .of_match_table = nbody_of_match,
351     },
352 };
353
354 module_platform_driver(nbody_driver);

```

```

355
356 MODULE_LICENSE("GPL");
357 MODULE_AUTHOR("CSEE4840 Group");
358 MODULE_DESCRIPTION("N-body Avalon-MM accelerator driver");

```

### A.1.11 nbody\_ioctl.h

Listing 13: code/Software/nbody\_ioctl.h

```

1  #ifndef _NBODY_IOCTL_H
2  #define _NBODY_IOCTL_H
3
4  #ifdef __KERNEL__
5  #include <linux/ioctl.h>
6  #include <linux/types.h>
7  #else
8  #include <sys/ioctl.h>
9  #include <stdint.h>
10 #endif
11
12 #define NBODY_MAX_BODIES 1024
13 #define NBODY_DATA_W 27
14 #define NBODY_DATA_MASK 0x07ffffffu
15 #define NBODY_F27_SIGN_SHIFT 26
16 #define NBODY_F27_EXP_SHIFT 18
17 #define NBODY_F27_MANT_MASK 0x0003ffffu
18 #define NBODY_MAGIC 'n'
19
20 typedef struct {
21     float x;
22     float y;
23     float mass;
24     float vx;
25     float vy;
26 } nbody_particle_t;
27
28 typedef struct {
29     float x;
30     float y;
31 } nbody_result_t;
32
33 typedef nbody_result_t body_pos_t;
34
35 typedef struct {
36     uint32_t num_bodies;
37     uint32_t gap;
38 } nbody_config_t;
39
40 typedef struct {
41     const nbody_particle_t *particles;
42     uint32_t count;
43 } nbody_bodies_arg_t;
44
45 typedef struct {
46     nbody_result_t *results;
47     uint32_t count;
48 } nbody_read_arg_t;
49
50 #define NBODY_WRITE_CONFIG    _IOW(NBODY_MAGIC, 1, nbody_config_t)
51 #define NBODY_WRITE_BODIES    _IOW(NBODY_MAGIC, 2, nbody_bodies_arg_t)
52 #define NBODY_START_RUN      _IO(NBODY_MAGIC, 3)
53 #define NBODY_CHECK_DONE     _IOR(NBODY_MAGIC, 4, int)
54 #define NBODY_READ_RESULTS   _IOWR(NBODY_MAGIC, 5, nbody_read_arg_t)
55 #define NBODY_CLEAR_READ     _IO(NBODY_MAGIC, 6)
56 #define NBODY_STOP           _IO(NBODY_MAGIC, 7)
57 #define NBODY_SOFT_RESET     _IO(NBODY_MAGIC, 8)
58
59 #define NBODY_WRITE_INIT_DATA NBODY_WRITE_BODIES
60 #define NBODY_READ_POSITIONS  NBODY_READ_RESULTS

```

```
61
62 #endif
```

### A.1.12 display\_driver.c

Listing 14: code/Software/display\_driver.c

```
1  /*
2  * Packed bitmap display Avalon-MM Linux driver.
3  *
4  * The kernel only transfers the userspace-rendered 640x480 1-bpp framebuffer
5  * to vga_bitmap_avmm.sv. It does not render particles or text.
6  */
7
8  #include <linux/errno.h>
9  #include <linux/fs.h>
10 #include <linux/init.h>
11 #include <linux/io.h>
12 #include <linux/kernel.h>
13 #include <linux/miscdevice.h>
14 #include <linux/module.h>
15 #include <linux/of.h>
16 #include <linux/of_address.h>
17 #include <linux/platform_device.h>
18 #include <linux/slab.h>
19 #include <linux/uaccess.h>
20 #include "display_ioctl.h"
21
22 #define DRIVER_NAME "nbody_display"
23
24 struct display_dev {
25     struct resource res;
26     void __iomem *virtbase;
27     uint32_t *kbuf;
28 };
29
30 static struct display_dev dev;
31
32 /*
33 * Main ioctl dispatch function for /dev/nbody_display.
34 *
35 * DISPLAY_WRITE_FRAME copies a full packed framebuffer from userspace and
36 * writes it into the hardware framebuffer. DISPLAY_CLEAR clears every hardware
37 * framebuffer word to zero.
38 */
39 static long display_ioctl(struct file *f, unsigned int cmd, unsigned long arg)
40 {
41     uint32_t i;
42
43     switch (cmd) {
44     case DISPLAY_WRITE_FRAME:
45         if (copy_from_user(dev.kbuf, (uint32_t __user *)arg,
46                             DISPLAY_WORDS * sizeof(uint32_t)))
47             return -EFAULT;
48
49         /*
50          * vga_bitmap_avmm.sv uses word addresses 0..9599, so byte offset
51          * is i * 4. Framebuffer packing is word = y * 20 + x / 32,
52          * bit = x % 32, LSB-first.
53          */
54         for (i = 0; i < DISPLAY_WORDS; i++)
55             iowrite32(dev.kbuf[i], dev.virtbase + i * 4);
56         return 0;
57
58     case DISPLAY_CLEAR:
59         for (i = 0; i < DISPLAY_WORDS; i++)
60             iowrite32(0, dev.virtbase + i * 4);
61         return 0;
62     }
```

```

63     default:
64         return -EINVAL;
65     }
66 }
67
68 static const struct file_operations display_fops = {
69     .owner = THIS_MODULE,
70     .unlocked_ioctl = display_ioctl,
71 };
72
73 static struct miscdevice display_misc_device = {
74     .minor = MISC_DYNAMIC_MINOR,
75     .name = "nbody_display",
76     .fops = &display_fops,
77 };
78
79 /*
80  * Platform driver probe function.
81  *
82  * This function is called when the matching display device-tree node is found.
83  * It maps the Avalon-MM framebuffer, allocates a kernel-side frame buffer, and
84  * registers /dev/nbody_display as a misc device.
85  */
86 static int display_probe(struct platform_device *pdev)
87 {
88     int ret;
89
90     ret = of_address_to_resource(pdev->dev.of_node, 0, &dev.res);
91     if (ret)
92         return ret;
93
94     if (!request_mem_region(dev.res.start, resource_size(&dev.res), DRIVER_NAME))
95         return -EBUSY;
96
97     dev.virtbase = of_iomap(pdev->dev.of_node, 0);
98     if (!dev.virtbase) {
99         ret = -ENOMEM;
100         goto out_release;
101     }
102
103     dev.kbuf = kmalloc(DISPLAY_WORDS * sizeof(uint32_t), GFP_KERNEL);
104     if (!dev.kbuf) {
105         ret = -ENOMEM;
106         goto out_unmap;
107     }
108
109     ret = misc_register(&display_misc_device);
110     if (ret)
111         goto out_free;
112
113     pr_info("display driver registered /dev/nbody_display\n");
114     return 0;
115
116 out_free:
117     kfree(dev.kbuf);
118 out_unmap:
119     iounmap(dev.virtbase);
120 out_release:
121     release_mem_region(dev.res.start, resource_size(&dev.res));
122     return ret;
123 }
124
125 /*
126  * Platform driver remove function.
127  *
128  * Release /dev/nbody_display, free the kernel frame buffer, unmap the hardware
129  * framebuffer, and release the memory region.
130  */
131 static int display_remove(struct platform_device *pdev)
132 {
133     misc_deregister(&display_misc_device);
134     kfree(dev.kbuf);
135     if (dev.virtbase)

```

```

136     iounmap(dev.virtbase);
137     release_mem_region(dev.res.start, resource_size(&dev.res));
138     return 0;
139 }
140
141 static const struct of_device_id display_of_match[] = {
142     { .compatible = "csee4840,vga_bitmap_avmm-1.0" },
143     {}
144 };
145 MODULE_DEVICE_TABLE(of, display_of_match);
146
147 static struct platform_driver display_driver = {
148     .probe = display_probe,
149     .remove = display_remove,
150     .driver = {
151         .name = DRIVER_NAME,
152         .of_match_table = display_of_match,
153     },
154 };
155
156 module_platform_driver(display_driver);
157
158 MODULE_LICENSE("GPL");
159 MODULE_AUTHOR("CSEE4840 Group");
160 MODULE_DESCRIPTION("Packed VGA bitmap Avalon-MM display driver");

```

### A.1.13 display\_ioctl.h

Listing 15: code/Software/display\_ioctl.h

```

1  #ifndef _DISPLAY_IOCTL_H
2  #define _DISPLAY_IOCTL_H
3
4  #ifdef __KERNEL__
5  #include <linux/ioctl.h>
6  #include <linux/types.h>
7  #else
8  #include <sys/ioctl.h>
9  #include <stdint.h>
10 #endif
11
12 #define DISPLAY_WIDTH 640
13 #define DISPLAY_HEIGHT 480
14 #define DISPLAY_WORDS_PER_ROW 20
15 #define DISPLAY_WORDS 9600
16
17 #define DISPLAY_MAGIC 'd'
18
19 #define DISPLAY_WRITE_FRAME _IOW(DISPLAY_MAGIC, 1, uint32_t *)
20 #define DISPLAY_CLEAR _IO(DISPLAY_MAGIC, 2)
21
22 #endif

```

### A.1.14 Makefile

Listing 16: code/Software/Makefile

```

1  ifneq ($(KERNELRELEASE),)
2
3  obj-m += accelerator_driver.o
4  obj-m += display_driver.o
5
6  else
7
8  CC ?= cc

```

```

9 KDIR ?= /usr/src/linux-headers-$(shell uname -r)
10 PWD := $(shell pwd)
11
12 CFLAGS += -Wall -Wextra -pthread
13 LDLIBS += -libusb-1.0 -pthread
14
15 APP = nbody_app
16 OBJECTS = main.o nbody.o display.o keyboard.o usbkeyboard.o
17
18 TARFILES = Makefile main.c nbody.c nbody.h nbody_ioctl.h \
19           display.h display.c display_ioctl.h \
20           keyboard.h keyboard.c \
21           usbkeyboard.h usbkeyboard.c accelerator_driver.c display_driver.c
22
23 .PHONY: all app modules modules-clean clean distclean
24
25 all: app
26
27 app: $(APP)
28
29 $(APP): $(OBJECTS)
30     $(CC) $(CFLAGS) -o $@ $(OBJECTS) $(LDLIBS)
31
32 main.o: main.c display.h keyboard.h nbody.h nbody_ioctl.h display_ioctl.h
33 nbody.o: nbody.c nbody.h nbody_ioctl.h
34 display.o: display.c display.h display_ioctl.h nbody.h nbody_ioctl.h
35 keyboard.o: keyboard.c keyboard.h usbkeyboard.h nbody.h nbody_ioctl.h
36 usbkeyboard.o: usbkeyboard.c usbkeyboard.h
37
38 modules:
39     $(MAKE) -C $(KDIR) SUBDIRS=$(PWD) modules
40
41 modules-clean:
42     $(MAKE) -C $(KDIR) SUBDIRS=$(PWD) clean
43
44 nbody_app.tar.gz: $(TARFILES)
45     rm -rf nbody_app_dir
46     mkdir nbody_app_dir
47     ln $(TARFILES) nbody_app_dir
48     tar zcf nbody_app.tar.gz nbody_app_dir
49     rm -rf nbody_app_dir
50
51 clean:
52     rm -rf $(OBJECTS) $(APP) nbody_app.tar.gz nbody_app_dir
53
54 distclean: clean modules-clean
55
56 endif

```

## A.2 Software Test Code

### A.2.1 avmm\_smoke\_test.c

Listing 17: code/Software/test/avmm\_smoke\_test.c

```

1 /*
2  * avmm_smoke_test.c
3  *
4  * Minimal Linux userspace smoke test for DE1-SoC HPS -> FPGA lightweight
5  * Avalon-MM peripherals:
6  *   - nbody_accel_avmm
7  *   - vga_bitmap_avmm
8  *
9  * Build on the board:
10 *   gcc -O2 -Wall -Wextra -o avmm_smoke_test avmm_smoke_test.c
11 *
12 * Run as root / with permission to open /dev/mem:
13 *   ./avmm_smoke_test <nbody_offset_hex> <vga_offset_hex>

```

```

14  */
15
16  #include <errno.h>
17  #include <fcntl.h>
18  #include <inttypes.h>
19  #include <stdbool.h>
20  #include <stdint.h>
21  #include <stdio.h>
22  #include <stdlib.h>
23  #include <string.h>
24  #include <sys/mman.h>
25  #include <unistd.h>
26
27  #define LW_BRIDGE_BASE 0xFF200000u
28  #define LW_BRIDGE_SPAN 0x00200000u
29
30  /* nbody_accel_avmm register word indices */
31  enum {
32      NB_GO          = 0x00,
33      NB_N_BODIES   = 0x01,
34      NB_GAP        = 0x02,
35      NB_X_IN       = 0x03,
36      NB_Y_IN       = 0x04,
37      NB_M_IN       = 0x05,
38      NB_VX_IN      = 0x06,
39      NB_VY_IN      = 0x07,
40      NB_DONE       = 0x08,
41      NB_READ       = 0x09,
42      NB_OUT_X      = 0x0A,
43      NB_OUT_Y      = 0x0B,
44  };
45
46  /* vga_bitmap_avmm constants */
47  #define VGA_WIDTH      640u
48  #define VGA_HEIGHT     480u
49  #define VGA_WORDS_PER_ROW 20u
50  #define VGA_FB_WORDS   (VGA_HEIGHT * VGA_WORDS_PER_ROW)
51
52  #define FP27_SIGN_SHIFT 26u
53  #define FP27_EXP_SHIFT  18u
54  #define FP27_MANT_MASK  0x0003FFFFu
55
56  static inline void mmio_write(volatile uint32_t *base, uint32_t word_index, uint32_t value)
57  {
58      base[word_index] = value;
59      __sync_synchronize();
60  }
61
62  static inline uint32_t mmio_read(volatile uint32_t *base, uint32_t word_index) {
63      uint32_t value = base[word_index];
64      __sync_synchronize();
65      return value;
66  }
67
68  static uint32_t f32_to_f27(float value) {
69      uint32_t bits;
70      uint32_t sign;
71      uint32_t exp;
72      uint32_t mant;
73      uint32_t rounded_mant;
74
75      memcpy(&bits, &value, sizeof(bits));
76      sign = (bits >> 31) & 1u;
77      exp = (bits >> 23) & 0xffu;
78      mant = bits & 0x007ffffu;
79
80      if (exp == 0u) {
81          return 0u;
82      }
83      if (exp == 0xffu) {
84          return (sign << FP27_SIGN_SHIFT) | (0xfeu << FP27_EXP_SHIFT) | FP27_MANT_MASK;
85      }

```

```

86     rounded_mant = (mant + 0x10u) >> 5;
87     if (rounded_mant == (1u << 18)) {
88         rounded_mant = 0u;
89         exp++;
90         if (exp >= 0xffu) {
91             exp = 0xfeu;
92             rounded_mant = FP27_MANT_MASK;
93         }
94     }
95
96     return (sign << FP27_SIGN_SHIFT) | (exp << FP27_EXP_SHIFT) |
97           (rounded_mant & FP27_MANT_MASK);
98 }
99
100 static uint32_t parse_hex_arg(const char *s, const char *name) {
101     char *end = NULL;
102     errno = 0;
103     unsigned long value = strtoul(s, &end, 0);
104     if (errno != 0 || end == s || *end != '\0' || value > 0xFFFFFFFFul) {
105         fprintf(stderr, "Bad %s: %s\n", name, s);
106         exit(2);
107     }
108     if ((value & 0x3u) != 0) {
109         fprintf(stderr, "Warning: %s is not 4-byte aligned: 0x%08lx\n", name, value);
110     }
111     return (uint32_t)value;
112 }
113
114 static bool poll_done(volatile uint32_t *nb, unsigned max_iters) {
115     for (unsigned i = 0; i < max_iters; ++i) {
116         if (mmio_read(nb, NB_DONE) & 0x1u) {
117             return true;
118         }
119     }
120     return false;
121 }
122
123 static void vga_clear(volatile uint32_t *vga, uint32_t value) {
124     for (uint32_t i = 0; i < VGA_FB_WORDS; ++i) {
125         mmio_write(vga, i, value);
126     }
127 }
128
129 static void vga_draw_test_pattern(volatile uint32_t *vga) {
130     /* Clear screen. */
131     vga_clear(vga, 0x00000000u);
132
133     /* Draw a simple border plus a diagonal-ish line. */
134     for (uint32_t y = 0; y < VGA_HEIGHT; ++y) {
135         for (uint32_t xword = 0; xword < VGA_WORDS_PER_ROW; ++xword) {
136             uint32_t word = 0;
137             uint32_t x0 = xword * 32u;
138             for (uint32_t b = 0; b < 32u; ++b) {
139                 uint32_t x = x0 + b;
140                 bool border = (x == 0u || x == VGA_WIDTH - 1u || y == 0u || y == VGA_HEIGHT
141 - 1u);
141                 bool diag = (x < VGA_WIDTH && (x / 2u) == y);
142                 if (border || diag) {
143                     word |= (1u << b); /* LSB-first, matching hardware comment. */
144                 }
145             }
146             mmio_write(vga, y * VGA_WORDS_PER_ROW + xword, word);
147         }
148     }
149 }
150
151 static int test_nbody_zero_body_done(volatile uint32_t *nb) {
152     printf("[nbody] zero-body GO/DONE/READ smoke test...\n");
153
154     mmio_write(nb, NB_READ, 0u);
155     mmio_write(nb, NB_N_BODIES, 0u);
156     mmio_write(nb, NB_GAP, 1u);
157     mmio_write(nb, NB_GO, 1u);

```

```

158
159     if (!poll_done(nb, 1000000u)) {
160         fprintf(stderr, "[FAIL] DONE did not go high for N_BODIES=0.\n");
161         return 1;
162     }
163     printf("[pass] DONE went high.\n");
164
165     mmio_write(nb, NB_READ, 0u);
166     for (unsigned i = 0; i < 1000u; ++i) {
167         if ((mmio_read(nb, NB_DONE) & 0x1u) == 0u) {
168             printf("[pass] DONE cleared after READ=0.\n");
169             return 0;
170         }
171     }
172
173     fprintf(stderr, "[FAIL] DONE did not clear after READ=0.\n");
174     return 1;
175 }
176
177 static int test_nbody_one_body_output(volatile uint32_t *nb) {
178     printf("[nbody] one-body input/output smoke test...\n");
179
180     /* Load one body: x=1.0, y=2.0, m=1.0, vx=0, vy=0. */
181     const uint32_t x_in = f32_to_f27(1.0f);
182     const uint32_t y_in = f32_to_f27(2.0f);
183     const uint32_t m_in = f32_to_f27(1.0f);
184     const uint32_t zero = f32_to_f27(0.0f);
185
186     mmio_write(nb, NB_READ, 0u);
187     mmio_write(nb, NB_N_BODIES, 1u);
188     mmio_write(nb, NB_GAP, 1u);
189     mmio_write(nb, NB_X_IN, x_in);
190     mmio_write(nb, NB_Y_IN, y_in);
191     mmio_write(nb, NB_M_IN, m_in);
192     mmio_write(nb, NB_VX_IN, zero);
193     mmio_write(nb, NB_VY_IN, zero); /* commits body 0 */
194     mmio_write(nb, NB_GO, 1u);
195
196     if (!poll_done(nb, 10000000u)) {
197         fprintf(stderr, "[FAIL] DONE did not go high for N_BODIES=1.\n");
198         return 1;
199     }
200
201     uint32_t out_x = mmio_read(nb, NB_OUT_X) & 0x07FFFFFFu;
202     uint32_t out_y = mmio_read(nb, NB_OUT_Y) & 0x07FFFFFFu; /* increments output pointer */
203
204     printf("[info] OUT_X = 0x%08" PRIx32 " , OUT_Y = 0x%08" PRIx32 "\n", out_x, out_y);
205     printf("[info] Expected for current N=1 zero-velocity smoke case is usually x=0x%08x, y
206 =0x%08x.\n",
207         x_in, y_in);
208
209     if (out_x == x_in && out_y == y_in) {
210         printf("[pass] one-body output matches input position.\n");
211     } else {
212         printf("[warn] one-body output did not exactly match. This may indicate either a
213 real issue or that the integrator/first-step path changed the value.\n");
214     }
215
216     mmio_write(nb, NB_READ, 0u);
217     return 0;
218 }
219
220 int main(int argc, char **argv) {
221     if (argc != 3) {
222         fprintf(stderr,
223             "Usage: %s <nbody_offset_hex> <vga_offset_hex>\n"
224             "Example: %s 0x00000 0x01000\n",
225             argv[0], argv[0]);
226         return 2;
227     }
228
229     uint32_t nbody_offset = parse_hex_arg(argv[1], "nbody offset");
230     uint32_t vga_offset = parse_hex_arg(argv[2], "vga offset");

```

```

229
230     if (nbody_offset >= LW_BRIDGE_SPAN || vga_offset >= LW_BRIDGE_SPAN) {
231         fprintf(stderr, "Offsets must be inside LW bridge span 0x%08x.\n", LW_BRIDGE_SPAN);
232         return 2;
233     }
234
235     int fd = open("/dev/mem", O_RDWR | O_SYNC);
236     if (fd < 0) {
237         perror("open /dev/mem");
238         return 1;
239     }
240
241     void *map = mmap(NULL, LW_BRIDGE_SPAN, PROT_READ | PROT_WRITE, MAP_SHARED, fd,
242                     LW_BRIDGE_BASE);
243     if (map == MAP_FAILED) {
244         perror("mmap lightweight bridge");
245         close(fd);
246         return 1;
247     }
248     volatile uint32_t *nbody = (volatile uint32_t *)((volatile uint8_t *)map + nbody_offset)
249     ;
250     volatile uint32_t *vga     = (volatile uint32_t *)((volatile uint8_t *)map + vga_offset);
251
252     printf("LW bridge base: 0x%08x\n", LW_BRIDGE_BASE);
253     printf("nbody offset  : 0x%08" PRIx32 "\n", nbody_offset);
254     printf("vga offset      : 0x%08" PRIx32 "\n", vga_offset);
255
256     int failures = 0;
257
258     /* VGA readback is intentionally always zero in current RTL. */
259     printf("[vga] readback check: word0=0x%08" PRIx32 " word123=0x%08" PRIx32 " (current RTL
260           should read 0)\n",
261           mmio_read(vga, 0), mmio_read(vga, 123));
262     printf("[vga] drawing border + diagonal test pattern...\n");
263     vga_draw_test_pattern(vga);
264     printf("[vga] done. Check monitor for white border and diagonal on black background.\n")
265     ;
266
267     failures += test_nbody_zero_body_done(nbody);
268     failures += test_nbody_one_body_output(nbody);
269
270     munmap(map, LW_BRIDGE_SPAN);
271     close(fd);
272
273     if (failures == 0) {
274         printf("All smoke tests completed.\n");
275         return 0;
276     }
277
278     fprintf(stderr, "%d smoke test(s) failed.\n", failures);
279     return 1;
280 }

```

## A.2.2 avmm\_frame\_xy\_dump.c

Listing 18: code/Software/test/avmm\_frame\_xy\_dump.c

```

1  #define _DEFAULT_SOURCE
2
3  /*
4  * avmm_frame_xy_dump.c
5  *
6  * Board-side N-body Avalon-MM test that loads a frame-input txt file, runs one
7  * accelerator pass, and writes the resulting X/Y positions to a txt file for
8  * comparison against a golden output.
9  *
10 * Input format matches the txt files in Hardware/tb/frame_input:
11 * # i px py vx vy m (S1E8M18 hex)

```

```

12 *      0 20CE1A0 2156136 0000000 0000000 1FA8587
13 *
14 * Build on the board:
15 * gcc -O2 -Wall -Wextra -std=c11 -o avmm_frame_xy_dump avmm_frame_xy_dump.c
16 *
17 * Run as root / with permission to open /dev/mem:
18 * ./avmm_frame_xy_dump <nbody_offset_hex> <input_frame.txt> <output_xy.txt> [n_bodies] [
    gap]
19 *
20 * Example:
21 * ./avmm_frame_xy_dump 0x00000 frame0_1024binit200_27bits.txt hw_xy_1024.txt 1024 1
22 */
23
24 #include <errno.h>
25 #include <fcntl.h>
26 #include <inttypes.h>
27 #include <stdbool.h>
28 #include <stdint.h>
29 #include <stdio.h>
30 #include <stdlib.h>
31 #include <string.h>
32 #include <sys/mman.h>
33 #include <unistd.h>
34
35 #define LW_BRIDGE_BASE 0xFF200000u
36 #define LW_BRIDGE_SPAN 0x00200000u
37
38 #define NBODY_MAX_BODIES 1024u
39 #define NBODY_DATA_MASK 0x07FFFFFFu
40
41 enum {
42     NB_GO          = 0x00,
43     NB_N_BODIES   = 0x01,
44     NB_GAP        = 0x02,
45     NB_X_IN       = 0x03,
46     NB_Y_IN       = 0x04,
47     NB_M_IN       = 0x05,
48     NB_VX_IN      = 0x06,
49     NB_VY_IN      = 0x07,
50     NB_DONE       = 0x08,
51     NB_READ       = 0x09,
52     NB_OUT_X      = 0x0A,
53     NB_OUT_Y      = 0x0B,
54 };
55
56 typedef struct {
57     uint32_t x;
58     uint32_t y;
59     uint32_t vx;
60     uint32_t vy;
61     uint32_t m;
62     bool present;
63 } frame_body_t;
64
65 static inline void mmio_write(volatile uint32_t *base, uint32_t word_index, uint32_t value)
66 {
67     base[word_index] = value;
68     __sync_synchronize();
69 }
70
71 static inline uint32_t mmio_read(volatile uint32_t *base, uint32_t word_index)
72 {
73     uint32_t value = base[word_index];
74     __sync_synchronize();
75     return value;
76 }
77
78 static uint32_t parse_u32_arg(const char *s, const char *name)
79 {
80     char *end = NULL;
81     errno = 0;
82     unsigned long value = strtoul(s, &end, 0);
83     if (errno != 0 || end == s || *end != '\0' || value > 0xFFFFFFFF) {

```

```

84     fprintf(stderr, "Bad %s: %s\n", name, s);
85     exit(2);
86 }
87 return (uint32_t)value;
88 }
89
90 static uint32_t parse_count_arg(const char *s, const char *name)
91 {
92     uint32_t value = parse_u32_arg(s, name);
93     if (value == 0u || value > NBODY_MAX_BODIES) {
94         fprintf(stderr, "%s must be in 1..%u, got %" PRIu32 "\n",
95                 name, NBODY_MAX_BODIES, value);
96         exit(2);
97     }
98     return value;
99 }
100
101 static uint32_t parse_gap_arg(const char *s)
102 {
103     uint32_t value = parse_u32_arg(s, "gap");
104     if (value == 0u) {
105         fprintf(stderr, "gap must be nonzero\n");
106         exit(2);
107     }
108     return value;
109 }
110
111 static uint32_t parse_offset_arg(const char *s)
112 {
113     uint32_t value = parse_u32_arg(s, "nbody offset");
114     if ((value & 0x3u) != 0u) {
115         fprintf(stderr, "Warning: nbody offset is not 4-byte aligned: 0x%08" PRIx32 "\n",
116                 value);
117     }
118     if (value >= LW_BRIDGE_SPAN) {
119         fprintf(stderr, "nbody offset must be inside LW bridge span 0x%08x\n",
120                 LW_BRIDGE_SPAN);
121         exit(2);
122     }
123     return value;
124 }
125
126 static uint32_t infer_count(const frame_body_t bodies[NBODY_MAX_BODIES])
127 {
128     uint32_t count = 0;
129
130     for (uint32_t i = 0; i < NBODY_MAX_BODIES; ++i) {
131         if (bodies[i].present) {
132             count = i + 1u;
133         }
134     }
135
136     return count;
137 }
138
139 static int read_frame_file(const char *path, frame_body_t bodies[NBODY_MAX_BODIES],
140                          uint32_t *inferred_count)
141 {
142     FILE *fp;
143     char line[512];
144     unsigned line_no = 0;
145
146     memset(bodies, 0, sizeof(frame_body_t) * NBODY_MAX_BODIES);
147
148     fp = fopen(path, "r");
149     if (!fp) {
150         perror(path);
151         return 1;
152     }
153
154     while (fgets(line, sizeof(line), fp)) {
155         int idx;
156         unsigned px;

```

```

157     unsigned py;
158     unsigned vx;
159     unsigned vy;
160     unsigned m;
161     char tail;
162     int got;
163     char *p = line;
164
165     line_no++;
166     while (*p == ' ' || *p == '\t') {
167         p++;
168     }
169     if (*p == '\0' || *p == '\n' || *p == '#') {
170         continue;
171     }
172
173     got = sscanf(p, "%d %x %x %x %x %c", &idx, &px, &py, &vx, &vy, &m, &tail);
174     if (got != 6) {
175         fprintf(stderr, "%s:%u: expected six columns: i px py vx vy m\n",
176             path, line_no);
177         fclose(fp);
178         return 1;
179     }
180     if (idx < 0 || idx >= (int)NBODY_MAX_BODIES) {
181         fprintf(stderr, "%s:%u: body index %d outside 0..%u\n",
182             path, line_no, idx, NBODY_MAX_BODIES - 1u);
183         fclose(fp);
184         return 1;
185     }
186     if ((px | py | vx | vy | m) & ~NBODY_DATA_MASK) {
187         fprintf(stderr, "%s:%u: S1E8M18 fields must fit in 27 bits\n",
188             path, line_no);
189         fclose(fp);
190         return 1;
191     }
192     if (bodies[idx].present) {
193         fprintf(stderr, "%s:%u: duplicate body index %d\n", path, line_no, idx);
194         fclose(fp);
195         return 1;
196     }
197
198     bodies[idx].x = px;
199     bodies[idx].y = py;
200     bodies[idx].vx = vx;
201     bodies[idx].vy = vy;
202     bodies[idx].m = m;
203     bodies[idx].present = true;
204 }
205
206 if (ferror(fp)) {
207     perror(path);
208     fclose(fp);
209     return 1;
210 }
211
212 fclose(fp);
213 *inferred_count = infer_count(bodies);
214 if (*inferred_count == 0u) {
215     fprintf(stderr, "%s: no body rows found\n", path);
216     return 1;
217 }
218
219 return 0;
220 }
221
222 static int validate_frame_prefix(const frame_body_t bodies[NBODY_MAX_BODIES],
223     uint32_t count)
224 {
225     for (uint32_t i = 0; i < count; ++i) {
226         if (!bodies[i].present) {
227             fprintf(stderr, "input frame is missing required body index %" PRIu32 "\n", i);
228             return 1;
229         }

```

```

230     }
231     return 0;
232 }
233
234 static void load_frame(volatile uint32_t *nb, const frame_body_t bodies[NBODY_MAX_BODIES],
235                      uint32_t count, uint32_t gap)
236 {
237     mmio_write(nb, NB_READ, 0u);
238     mmio_write(nb, NB_N_BODIES, count);
239     mmio_write(nb, NB_GAP, gap);
240
241     for (uint32_t i = 0; i < count; ++i) {
242         mmio_write(nb, NB_X_IN, bodies[i].x);
243         mmio_write(nb, NB_Y_IN, bodies[i].y);
244         mmio_write(nb, NB_M_IN, bodies[i].m);
245         mmio_write(nb, NB_VX_IN, bodies[i].vx);
246         mmio_write(nb, NB_VY_IN, bodies[i].vy);
247     }
248 }
249
250 static bool poll_done(volatile uint32_t *nb, uint32_t max_iters)
251 {
252     for (uint32_t i = 0; i < max_iters; ++i) {
253         if (mmio_read(nb, NB_DONE) & 0x1u) {
254             return true;
255         }
256         if ((i & 0x3ffu) == 0u) {
257             usleep(1000);
258         }
259     }
260     return false;
261 }
262
263 static int run_and_write_xy(volatile uint32_t *nb, const char *out_path,
264                            const char *input_path, uint32_t count, uint32_t gap)
265 {
266     FILE *out;
267
268     mmio_write(nb, NB_GO, 1u);
269     if (!poll_done(nb, 20000000u)) {
270         fprintf(stderr, "[FAIL] DONE did not go high for count=%" PRIu32 " , gap=%" PRIu32 "\n",
271                count, gap);
272         return 1;
273     }
274
275     out = fopen(out_path, "w");
276     if (!out) {
277         perror(out_path);
278         return 1;
279     }
280
281     fprintf(out, "# source %s\n", input_path);
282     fprintf(out, "# n_bodies %" PRIu32 "\n", count);
283     fprintf(out, "# gap %" PRIu32 "\n", gap);
284     fprintf(out, "# i x y (S1E8M18 hex)\n");
285
286     mmio_write(nb, NB_READ, 1u);
287     for (uint32_t i = 0; i < count; ++i) {
288         uint32_t x = mmio_read(nb, NB_OUT_X) & NBODY_DATA_MASK;
289         uint32_t y = mmio_read(nb, NB_OUT_Y) & NBODY_DATA_MASK;
290         fprintf(out, "%4" PRIu32 " %07" PRIx32 " %07" PRIx32 "\n", i, x, y);
291     }
292     mmio_write(nb, NB_READ, 0u);
293
294     if (fclose(out) != 0) {
295         perror(out_path);
296         return 1;
297     }
298
299     return 0;
300 }
301

```

```

302 int main(int argc, char **argv)
303 {
304     frame_body_t bodies[NBODY_MAX_BODIES];
305     uint32_t nbody_offset;
306     uint32_t inferred_count;
307     uint32_t count;
308     uint32_t gap;
309     int fd;
310     void *map;
311     volatile uint32_t *nbody;
312     int ret;
313
314     if (argc < 4 || argc > 6) {
315         fprintf(stderr,
316             "Usage: %s <nbody_offset_hex> <input_frame.txt> <output_xy.txt> [n_bodies] [
gap]\n"
317             "Example: %s 0x00000 frame0_1024binit200_27bits.txt hw_xy_1024.txt 1024 1\n"
318             ,
319             argv[0], argv[0]);
320     }
321
322     nbody_offset = parse_offset_arg(argv[1]);
323     gap = (argc >= 6) ? parse_gap_arg(argv[5]) : 1u;
324
325     if (read_frame_file(argv[2], bodies, &inferred_count) != 0) {
326         return 1;
327     }
328
329     count = (argc >= 5) ? parse_count_arg(argv[4], "n_bodies") : inferred_count;
330     if (count > inferred_count) {
331         fprintf(stderr,
332             "requested n_bodies=%" PRIu32 " but frame only contains rows through index %"
333             " PRIu32 "\n",
334             count, inferred_count - 1u);
335         return 2;
336     }
337     if (validate_frame_prefix(bodies, count) != 0) {
338         return 1;
339     }
340
341     fd = open("/dev/mem", O_RDWR | O_SYNC);
342     if (fd < 0) {
343         perror("open /dev/mem");
344         return 1;
345     }
346
347     map = mmap(NULL, LW_BRIDGE_SPAN, PROT_READ | PROT_WRITE, MAP_SHARED, fd, LW_BRIDGE_BASE)
;
348     if (map == MAP_FAILED) {
349         perror("mmap lightweight bridge");
350         close(fd);
351         return 1;
352     }
353
354     nbody = (volatile uint32_t *)((volatile uint8_t *)map + nbody_offset);
355
356     printf("LW bridge base: 0x%08x\n", LW_BRIDGE_BASE);
357     printf("nbody offset : 0x%08" PRIx32 "\n", nbody_offset);
358     printf("input frame : %s\n", argv[2]);
359     printf("output xy : %s\n", argv[3]);
360     printf("n_bodies : %" PRIu32 "\n", count);
361     printf("gap : %" PRIu32 "\n", gap);
362     printf("[nbody] loading frame...\n");
363
364     load_frame(nbody, bodies, count, gap);
365     printf("[nbody] running accelerator and dumping X/Y results...\n");
366     ret = run_and_write_xy(nbody, argv[3], argv[2], count, gap);
367
368     munmap(map, LW_BRIDGE_SPAN);
369     close(fd);
370
371     if (ret == 0) {

```

```

371     printf("[pass] wrote %s\n", argv[3]);
372 }
373 return ret;
374 }

```

## A.2.3 Test Makefile

Listing 19: code/Software/test/Makefile

```

1 # Makefile for direct userspace Avalon-MM smoke test on DE1-SoC
2 #
3 # This is intentionally simpler than the Lab 3 driver Makefile:
4 # it only builds a userspace program that opens /dev/mem and writes/reads
5 # the FPGA lightweight bridge directly. No .ko driver is built here.
6 #
7 # Usage on DE1-SoC:
8 #   make
9 #   make run NBODY_OFFSET=0x00000 VGA_OFFSET=0x01000
10 #   make run-frame NBODY_OFFSET=0x00000 FRAME_INPUT=frame0_1024binit200_27bits.txt
11 #       FRAME_OUTPUT=hw_xy_1024.txt
12 #
13 # Replace the offsets with Platform Designer assigned offsets
14 # relative to the HPS lightweight bridge base 0xFF200000.
15 CC      ?= cc
16 CFLAGS  ?= -O2 -Wall -Wextra -std=c11
17 SMOKE_TARGET := avmm_smoke_test
18 FRAME_TARGET := avmm_frame_xy_dump
19 TARGETS    := $(SMOKE_TARGET) $(FRAME_TARGET)
20
21 # Override these at the command line:
22 #   make run NBODY_OFFSET=0x00000 VGA_OFFSET=0x01000
23 NBODY_OFFSET ?= 0x00000
24 VGA_OFFSET   ?= 0x01000
25 FRAME_INPUT  ?= ../../Hardware/tb/frame_input/frame0_1024binit200_27bits.txt
26 FRAME_OUTPUT ?= hw_xy_1024.txt
27 FRAME_N      ?= 1024
28 FRAME_GAP    ?= 1
29
30 .PHONY: all run clean help
31
32 all: $(TARGETS)
33
34 $(SMOKE_TARGET): avmm_smoke_test.c
35     $(CC) $(CFLAGS) -o $@ $^
36
37 $(FRAME_TARGET): avmm_frame_xy_dump.c
38     $(CC) $(CFLAGS) -o $@ $^
39
40 run: $(SMOKE_TARGET)
41     ./$(SMOKE_TARGET) $(NBODY_OFFSET) $(VGA_OFFSET)
42
43 run-frame: $(FRAME_TARGET)
44     ./$(FRAME_TARGET) $(NBODY_OFFSET) $(FRAME_INPUT) $(FRAME_OUTPUT) $(FRAME_N) $(FRAME_GAP)
45
46 clean:
47     $(RM) $(TARGETS) *.o
48
49 help:
50     @echo "Build: make"
51     @echo "Run:   make run NBODY_OFFSET=0x<nbody_offset> VGA_OFFSET=0x<vga_offset>"
52     @echo "Example: make run NBODY_OFFSET=0x00000 VGA_OFFSET=0x01000"
53     @echo "Frame: make run-frame NBODY_OFFSET=0x<nbody_offset> FRAME_INPUT=<frame.txt>
54     FRAME_OUTPUT=<xy.txt> FRAME_N=<n> FRAME_GAP=<gap>"
55     @echo "Example: make run-frame NBODY_OFFSET=0x00000 FRAME_INPUT=../../Hardware/tb/
56     frame_input/frame0_1024binit200_27bits.txt FRAME_OUTPUT=hw_xy_1024.txt FRAME_N=1024
57     FRAME_GAP=1"

```

## A.3 Hardware Source Code

### A.3.1 nbody\_accel\_avmm.sv

Listing 20: code/Hardware/code/nbody\_accel\_avmm.sv

```
1 // Avalon-MM top-level shell for the N-body accelerator.
2 //
3 // Register map, 32-bit words:
4 //   0x00 GO           W Write bit 0 high to start one accelerator run. This also
5 //                   resets input/output body pointers to 0.
6 //   0x01 N_BODIES    W Number of active bodies in the simulation.
7 //   0x02 GAP         W Number of timesteps executed internally between DONE pulses.
8 //   0x03 X_IN       W Input X position for the current body.
9 //   0x04 Y_IN       W Input Y position for the current body.
10 //   0x05 M_IN       W Input mass for the current body.
11 //   0x06 VX_IN      W Input X velocity for the current body.
12 //   0x07 VY_IN      W Input Y velocity for the current body. Writing this register
13 //                   commits the current body and increments the input pointer.
14 //   0x08 DONE       R High when GAP timesteps have completed.
15 //   0x09 READ       W Write 1 while software is reading DONE/output data. Write
16 //                   0 after reading all outputs; this clears DONE and arms the
17 //                   next GO.
18 //   0x0A OUT_X      R Output X position for the current output body.
19 //   0x0B OUT_Y      R Output Y position for the current output body. Reading this
20 //                   register increments the output pointer.
21 // Data uses the lower 27 bits of each 32-bit Avalon word.
22
23 module nbody_accel_avmm #(
24     parameter int MAX_BODIES = 1024
25 ) (
26     input logic clk,
27     input logic reset, // active-high Platform Designer reset
28
29     input logic chipselect,
30     input logic read,
31     input logic write,
32     input logic [7:0] address,
33     input logic [31:0] writedata,
34     output logic [31:0] readdata
35 );
36
37 localparam logic [7:0] REG_GO = 8'h00;
38 localparam logic [7:0] REG_N_BODIES = 8'h01;
39 localparam logic [7:0] REG_GAP = 8'h02;
40 localparam logic [7:0] REG_X_IN = 8'h03;
41 localparam logic [7:0] REG_Y_IN = 8'h04;
42 localparam logic [7:0] REG_M_IN = 8'h05;
43 localparam logic [7:0] REG_VX_IN = 8'h06;
44 localparam logic [7:0] REG_VY_IN = 8'h07;
45 localparam logic [7:0] REG_DONE = 8'h08;
46 localparam logic [7:0] REG_READ = 8'h09;
47 localparam logic [7:0] REG_OUT_X = 8'h0A;
48 localparam logic [7:0] REG_OUT_Y = 8'h0B;
49
50 localparam int DATA_W = 27;
51 localparam int PAD_W = 32 - DATA_W;
52 localparam int PTR_W = $clog2(MAX_BODIES);
53 localparam logic [PTR_W-1:0] MAX_BODY_PTR = PTR_W'(MAX_BODIES - 1);
54
55 logic go_pulse;
56 logic read_reg;
57 logic out_y_read_d;
58 logic done;
59 logic first_step_pending;
60 logic first_step_for_run;
61
62 logic [31:0] n_bodies_reg;
63 logic [31:0] gap_reg;
64
65 logic [PTR_W-1:0] input_ptr;
```

```

66 logic [PTR_W-1:0] output_ptr;
67
68 logic [DATA_W-1:0] x_in_shadow;
69 logic [DATA_W-1:0] y_in_shadow;
70 logic [DATA_W-1:0] m_in_shadow;
71 logic [DATA_W-1:0] vx_in_shadow;
72 logic [DATA_W-1:0] vy_in_shadow;
73
74 logic          cpu_body_we;
75 logic [PTR_W-1:0] cpu_body_waddr;
76 logic [PTR_W-1:0] control_body_raddr;
77 logic [PTR_W-1:0] mem_body_raddr;
78 logic [DATA_W-1:0] body_x;
79 logic [DATA_W-1:0] body_y;
80 logic [DATA_W-1:0] body_m;
81 logic [DATA_W-1:0] body_vx;
82 logic [DATA_W-1:0] body_vy;
83 logic [DATA_W-1:0] body_ax;
84 logic [DATA_W-1:0] body_ay;
85
86 logic          body_update_we;
87 logic [PTR_W-1:0] body_update_addr;
88 logic [DATA_W-1:0] body_update_x;
89 logic [DATA_W-1:0] body_update_y;
90 logic [DATA_W-1:0] body_update_vx;
91 logic [DATA_W-1:0] body_update_vy;
92
93 logic          accel_we;
94 logic [PTR_W-1:0] accel_waddr;
95 logic [DATA_W-1:0] accel_ax;
96 logic [DATA_W-1:0] accel_ay;
97
98 assign mem_body_raddr = done ? output_ptr : control_body_raddr;
99
100 nbody_mem #(
101     .MAX_BODIES(MAX_BODIES),
102     .DATA_W(DATA_W),
103     .PTR_W(PTR_W)
104 ) u_mem (
105     .clk          (clk),
106
107     .cpu_body_we  (cpu_body_we),
108     .cpu_body_waddr (cpu_body_waddr),
109     .cpu_x        (x_in_shadow),
110     .cpu_y        (y_in_shadow),
111     .cpu_m        (m_in_shadow),
112     .cpu_vx       (vx_in_shadow),
113     .cpu_vy       (vy_in_shadow),
114
115     .body_raddr   (mem_body_raddr),
116     .body_x       (body_x),
117     .body_y       (body_y),
118     .body_m       (body_m),
119     .body_vx      (body_vx),
120     .body_vy      (body_vy),
121     .body_ax      (body_ax),
122     .body_ay      (body_ay),
123
124     .body_update_we (body_update_we),
125     .body_update_addr (body_update_addr),
126     .body_update_x  (body_update_x),
127     .body_update_y  (body_update_y),
128     .body_update_vx (body_update_vx),
129     .body_update_vy (body_update_vy),
130
131     .accel_we       (accel_we),
132     .accel_waddr    (accel_waddr),
133     .accel_ax       (accel_ax),
134     .accel_ay       (accel_ay)
135 );
136
137 nbody_control #(
138     .MAX_BODIES(MAX_BODIES),

```

```

139     .DATA_W(DATA_W)
140 ) u_control (
141     .clk          (clk),
142     .reset       (reset),
143
144     .go          (go_pulse),
145     .read_enable (read_reg),
146     .first_step  (first_step_for_run),
147     .n_bodies   (n_bodies_reg),
148     .gap        (gap_reg),
149     .done       (done),
150
151     .body_raddr  (control_body_raddr),
152     .body_x      (body_x),
153     .body_y      (body_y),
154     .body_m      (body_m),
155     .body_vx     (body_vx),
156     .body_vy     (body_vy),
157     .body_ax     (body_ax),
158     .body_ay     (body_ay),
159
160     .body_update_we (body_update_we),
161     .body_update_addr (body_update_addr),
162     .body_update_x  (body_update_x),
163     .body_update_y  (body_update_y),
164     .body_update_vx (body_update_vx),
165     .body_update_vy (body_update_vy),
166
167     .accel_we      (accel_we),
168     .accel_waddr   (accel_waddr),
169     .accel_ax     (accel_ax),
170     .accel_ay     (accel_ay)
171 );
172
173 always_ff @(posedge clk) begin
174     if (reset) begin
175         go_pulse      <= 1'b0;
176         read_reg      <= 1'b0;
177         out_y_read_d  <= 1'b0;
178         first_step_pending <= 1'b1;
179         first_step_for_run <= 1'b1;
180         n_bodies_reg <= 32'd0;
181         gap_reg       <= 32'd0;
182         input_ptr     <= '0;
183         output_ptr    <= '0;
184         x_in_shadow   <= '0;
185         y_in_shadow   <= '0;
186         m_in_shadow   <= '0;
187         vx_in_shadow  <= '0;
188         vy_in_shadow  <= '0;
189         cpu_body_we   <= 1'b0;
190         cpu_body_waddr <= '0;
191     end else begin
192         go_pulse      <= 1'b0;
193         cpu_body_we   <= 1'b0;
194         out_y_read_d <= chipselect && read && (address == REG_OUT_Y);
195
196         if (chipselect && write) begin
197             unique case (address)
198                 REG_GO: begin
199                     if (writedata[0]) begin
200                         go_pulse      <= 1'b1;
201                         read_reg      <= 1'b1;
202                         first_step_for_run <= first_step_pending;
203                         first_step_pending <= 1'b0;
204                         input_ptr     <= '0;
205                         output_ptr    <= '0;
206                     end
207                 end
208
209                 REG_N_BODIES: n_bodies_reg <= writedata;
210                 REG_GAP:      gap_reg      <= writedata;
211                 REG_X_IN:     x_in_shadow  <= writedata[DATA_W-1:0];

```

```

212     REG_Y_IN:    y_in_shadow  <= writedata[DATA_W-1:0];
213     REG_M_IN:    m_in_shadow  <= writedata[DATA_W-1:0];
214     REG_VX_IN:    vx_in_shadow <= writedata[DATA_W-1:0];
215
216     REG_VY_IN: begin
217         vy_in_shadow <= writedata[DATA_W-1:0];
218         cpu_body_we   <= 1'b1;
219         cpu_body_waddr <= input_ptr;
220         first_step_pending <= 1'b1;
221
222         if (input_ptr != MAX_BODY_PTR) begin
223             input_ptr <= input_ptr + 1'b1;
224         end
225     end
226
227     REG_READ: begin
228         read_reg <= writedata[0];
229         if (!writedata[0]) begin
230             output_ptr <= '0;
231         end
232     end
233
234     default: begin
235     end
236 endcase
237 end
238
239     if ((chipselct && read && (address == REG_OUT_Y)) && !out_y_read_d) begin
240         if (output_ptr != MAX_BODY_PTR) begin
241             output_ptr <= output_ptr + 1'b1;
242         end
243     end
244 end
245 end
246
247 always_comb begin
248     unique case (address)
249         REG_DONE: readdata = {31'd0, done};
250         REG_OUT_X: readdata = {{PAD_W{1'b0}}, body_x};
251         REG_OUT_Y: readdata = {{PAD_W{1'b0}}, body_y};
252         default: readdata = 32'd0;
253     endcase
254 end
255
256 endmodule

```

### A.3.2 vga\_bitmap\_avmm.sv

Listing 21: code/Hardware/code/vga\_bitmap\_avmm.sv

```

1  /*
2  * Avalon memory-mapped VGA bitmap peripheral.
3  *
4  * This is a separate Platform Designer/Qsys component from the n-body
5  * accelerator. Software writes a packed 640x480 1-bpp framebuffer through
6  * Avalon-MM; hardware only scans the bitmap out to VGA.
7  *
8  * Register map, 32-bit word addresses:
9  *   0..9599 framebuffer words, write-only from software
10 *   all reads from software return 0
11 *
12 * Bitmap packing:
13 *   WIDTH = 640, HEIGHT = 480, WORDS_PER_ROW = 20, FB_WORDS = 9600
14 *   word_index = y * 20 + x / 32
15 *   bit_index  = x % 32
16 *   frame[y * 20 + x / 32][x % 32] = pixel(x, y)
17 *
18 * Pixel value 0 displays black. Pixel value 1 displays white.
19 */

```

```

20
21 module vga_bitmap_avmm (
22     input  logic      clk,
23     input  logic      reset,          // active-high Platform Designer reset
24
25     input  logic      chipselect,
26     input  logic      read,
27     input  logic      write,
28     input  logic [13:0] address,      // word address, valid framebuffer range 0..9599
29     input  logic [31:0] writedata,
30     output logic [31:0] readdata,
31
32     output logic [7:0]  VGA_R,
33     output logic [7:0]  VGA_G,
34     output logic [7:0]  VGA_B,
35     output logic      VGA_CLK,
36     output logic      VGA_HS,
37     output logic      VGA_VS,
38     output logic      VGA_BLANK_n,
39     output logic      VGA_SYNC_n
40 );
41
42 localparam int WIDTH          = 640;
43 localparam int HEIGHT        = 480;
44 localparam int WORDS_PER_ROW = 20;
45 localparam int FB_WORDS      = HEIGHT * WORDS_PER_ROW;
46 localparam int ADDR_W        = 14;
47
48 localparam logic [ADDR_W-1:0] FB_LAST_WORD = ADDR_W'(FB_WORDS - 1);
49
50 logic [10:0] hcount;
51 logic [9:0]  vcount;
52
53 logic      av_addr_valid;
54 logic [13:0] av_ram_addr;
55
56 logic      vga_blank_raw;
57 logic      active_video;
58 logic      active_video_d;
59 logic [9:0] pixel_x;
60 logic [8:0] pixel_y;
61 logic [13:0] pixel_y_ext;
62 logic [13:0] scan_addr;
63 logic [4:0] bit_index;
64 logic [4:0] bit_index_d;
65 logic [31:0] scan_readdata;
66 logic      pixel_bit;
67
68 assign av_addr_valid = (address <= FB_LAST_WORD);
69 assign av_ram_addr   = av_addr_valid ? address : 14'd0;
70
71 // Reuse the Lab 3 VGA timing generator from vga_ball.sv.
72 // hcount[10:1] is the 640-pixel column and vcount is the 480-pixel row.
73 vga_counters counters (
74     .clk50      (clk),
75     .reset      (reset),
76     .hcount     (hcount),
77     .vcount     (vcount),
78     .VGA_CLK   (VGA_CLK),
79     .VGA_HS    (VGA_HS),
80     .VGA_VS    (VGA_VS),
81     .VGA_BLANK_n(vga_blank_raw),
82     .VGA_SYNC_n (VGA_SYNC_n)
83 );
84
85 assign active_video = vga_blank_raw;
86 assign pixel_x      = hcount[10:1];
87 assign pixel_y      = vcount[8:0];
88 assign pixel_y_ext  = {5'd0, pixel_y};
89
90 // word_index = y * 20 + x / 32 = (y << 4) + (y << 2) + x[9:5]
91 assign scan_addr = active_video ? ((pixel_y_ext << 4) + (pixel_y_ext << 2) +
92     {9'd0, pixel_x[9:5]}) :

```

```

93                                     14'd0;
94     assign bit_index = pixel_x[4:0];
95
96     framebuffer_ram #(
97         .FB_WORDS(FB_WORDS),
98         .ADDR_W (ADDR_W)
99     ) u_framebuffer (
100         .clk          (clk),
101
102         .port_a_we     (chipselct && write && av_addr_valid),
103         .port_a_addr   (av_ram_addr),
104         .port_a_writedata(writedata),
105
106         .port_b_addr   (scan_addr),
107         .port_b_readdata (scan_readdata)
108     );
109
110     always_ff @(posedge clk) begin
111         if (reset) begin
112             active_video_d <= 1'b0;
113             bit_index_d <= 5'd0;
114         end else begin
115             active_video_d <= active_video;
116             bit_index_d <= bit_index;
117         end
118     end
119
120     assign readdata = 32'd0;
121     assign VGA_BLANK_n = vga_blank_raw;
122     assign pixel_bit = scan_readdata[bit_index_d];
123
124     always_comb begin
125         if (active_video_d && pixel_bit) begin
126             VGA_R = 8'hff;
127             VGA_G = 8'hff;
128             VGA_B = 8'hff;
129         end else begin
130             VGA_R = 8'h00;
131             VGA_G = 8'h00;
132             VGA_B = 8'h00;
133         end
134     end
135
136 endmodule
137
138 module vga_counters(
139     input logic      clk50, reset,
140     output logic [10:0] hcount, // hcount[10:1] is pixel column
141     output logic [9:0] vcount, // vcount[9:0] is pixel row
142     output logic     VGA_CLK, VGA_HS, VGA_VS, VGA_BLANK_n, VGA_SYNC_n);
143
144 /*
145 * 640 X 480 VGA timing for a 50 MHz clock: one pixel every other cycle
146 *
147 * HCOUNT 1599 0           1279           1599 0
148 *
149 * -----|-----|-----|-----|
150 * |-----| Video |-----| Video
151 *
152 * |SYNC| BP |<-- HACTIVE -->|FP|SYNC| BP |<-- HACTIVE
153 *
154 * |----|-----|-----|-----|
155 * |----| VGA_HS |----|
156 */
157 // Parameters for hcount
158 parameter HACTIVE = 11'd 1280,
159           HFRONT_PORCH = 11'd 32,
160           HSYNC = 11'd 192,
161           HBACK_PORCH = 11'd 96,
162           HTOTAL = HACTIVE + HFRONT_PORCH + HSYNC +
163                 HBACK_PORCH; // 1600
164
165 // Parameters for vcount
166 parameter VACTIVE = 10'd 480,

```

```

166         VFRONT_PORCH = 10'd 10,
167         VSYNC        = 10'd 2,
168         VBACK_PORCH  = 10'd 33,
169         VTOTAL       = VACTIVE + VFRONT_PORCH + VSYNC +
170                     VBACK_PORCH; // 525
171
172     logic endOfLine;
173
174     always_ff @(posedge clk50)
175         if (reset)          hcount <= 0;
176         else if (endOfLine) hcount <= 0;
177         else                hcount <= hcount + 11'd 1;
178
179     assign endOfLine = hcount == HTOTAL - 1;
180
181     logic endOfField;
182
183     always_ff @(posedge clk50)
184         if (reset)          vcount <= 0;
185         else if (endOfLine)
186             if (endOfField) vcount <= 0;
187             else            vcount <= vcount + 10'd 1;
188
189     assign endOfField = vcount == VTOTAL - 1;
190
191     // Horizontal sync: from 0x520 to 0x5DF (0x57F)
192     // 101 0010 0000 to 101 1101 1111
193     assign VGA_HS = !( (hcount[10:8] == 3'b101) &
194                      !(hcount[7:5] == 3'b111));
195     assign VGA_VS = !( vcount[9:1] == (VACTIVE + VFRONT_PORCH) / 2);
196
197     assign VGA_SYNC_n = 1'b0; // For putting sync on the green signal; unused
198
199     // Horizontal active: 0 to 1279          Vertical active: 0 to 479
200     // 101 0000 0000 1280                  01 1110 0000 480
201     // 110 0011 1111 1599                  10 0000 1100 524
202     assign VGA_BLANK_n = !( hcount[10] & (hcount[9] | hcount[8]) ) &
203                          !( vcount[9] | (vcount[8:5] == 4'b1111) );
204
205     /* VGA_CLK is 25 MHz
206     *
207     * clk50    __|  |__|  |__|  |__|  |__|
208     *
209     * hcount[0]  -----|  -----|  -----|
210     */
211     assign VGA_CLK = hcount[0]; // 25 MHz clock: rising edge sensitive
212
213
214 endmodule

```

### A.3.3 nbody\_control.sv

Listing 22: code/Hardware/code/nbody\_control.sv

```

1  module nbody_control #(
2      parameter int MAX_BODIES = 1024,
3      parameter int DATA_W = 27
4  ) (
5      input  logic clk,
6      input  logic reset,
7
8      input  logic go,
9      input  logic read_enable,
10     input  logic first_step,
11     input  logic [31:0] n_bodies,
12     input  logic [31:0] gap,
13     output logic done,
14
15     output logic [$clog2(MAX_BODIES)-1:0] body_raddr,

```

```

16     input  logic [DATA_W-1:0]          body_x ,
17     input  logic [DATA_W-1:0]          body_y ,
18     input  logic [DATA_W-1:0]          body_m ,
19     input  logic [DATA_W-1:0]          body_vx ,
20     input  logic [DATA_W-1:0]          body_vy ,
21     input  logic [DATA_W-1:0]          body_ax ,
22     input  logic [DATA_W-1:0]          body_ay ,
23
24     output logic                        body_update_we ,
25     output logic [$clog2(MAX_BODIES)-1:0] body_update_addr ,
26     output logic [DATA_W-1:0]          body_update_x ,
27     output logic [DATA_W-1:0]          body_update_y ,
28     output logic [DATA_W-1:0]          body_update_vx ,
29     output logic [DATA_W-1:0]          body_update_vy ,
30
31     output logic                        accel_we ,
32     output logic [$clog2(MAX_BODIES)-1:0] accel_waddr ,
33     output logic [DATA_W-1:0]          accel_ax ,
34     output logic [DATA_W-1:0]          accel_ay
35 );
36
37     typedef enum logic [3:0] {
38         ST_IDLE ,
39         ST_LOAD_TILE_PRIME ,
40         ST_LOAD_TILE ,
41         ST_CLEAR_GROUP ,
42         ST_COMPUTE_PRIME ,
43         ST_COMPUTE_GROUP ,
44         ST_DRAIN_GROUP ,
45         ST_STORE_GROUP ,
46         ST_NEXT_GROUP ,
47         ST_NEXT_TILE ,
48         ST_INTEGRATE_PRIME ,
49         ST_INTEGRATE_START ,
50         ST_INTEGRATE_WAIT ,
51         ST_UPDATE_COMMIT ,
52         ST_DONE
53     } state_t;
54
55     typedef enum logic [1:0] {
56         UPD_NEXT_BODY ,
57         UPD_NEXT_STEP ,
58         UPD_DONE
59     } update_next_t;
60
61     localparam int PTR_W = $clog2(MAX_BODIES);
62     localparam logic [PTR_W-1:0] MAX_BODY_PTR = PTR_W'(MAX_BODIES - 1);
63     localparam logic [PTR_W-1:0] TILE_STRIDE = PTR_W'(16);
64     localparam int ACTIVE_W = PTR_W + 1;
65     localparam logic [31:0] MAX_BODIES_U32 = 32'(MAX_BODIES);
66     localparam int PIPE_LAT = 18;
67     localparam int WAIT_W = $clog2(PIPE_LAT + 4);
68
69     state_t state;
70     update_next_t update_next;
71
72     logic [31:0] timestep_count;
73     logic run_initial_half_step;
74     logic [PTR_W-1:0] tile_base;
75     logic [PTR_W-1:0] j_body_idx;
76     logic [PTR_W-1:0] current_j_idx;
77     logic [PTR_W-1:0] integrate_idx;
78     logic [WAIT_W-1:0] wait_count;
79     logic [1:0] compute_grp;
80     logic [1:0] store_lane;
81     logic [ACTIVE_W-1:0] active_count;
82     logic [ACTIVE_W-1:0] active_count_next;
83     logic core_rst_n;
84
85     logic core_clear_prev;
86     logic core_load_en;
87     logic core_compute_en;
88     logic [3:0] load_idx_count;

```

```

89 logic [3:0] core_load_idx;
90 logic [DATA_W-1:0] core_load_x;
91 logic [DATA_W-1:0] core_load_y;
92 logic [1:0] core_grp_sel;
93 logic [DATA_W-1:0] core_j_x;
94 logic [DATA_W-1:0] core_j_y;
95 logic [DATA_W-1:0] core_j_m;
96 logic [3:0] core_lane_mask;
97
98 logic [DATA_W-1:0] current_j_x;
99 logic [DATA_W-1:0] current_j_y;
100 logic [DATA_W-1:0] current_j_m;
101
102 logic [DATA_W-1:0] core_res_x [4];
103 logic [DATA_W-1:0] core_res_y [4];
104 logic unused_core_res_vld;
105
106 logic integrator_start;
107 logic integrator_done;
108 logic integrator_half_step;
109 logic [7:0] integrator_ax_half_exp;
110 logic [7:0] integrator_ay_half_exp;
111 logic [DATA_W-1:0] integrator_ax_in;
112 logic [DATA_W-1:0] integrator_ay_in;
113 logic [DATA_W-1:0] integrator_x_out;
114 logic [DATA_W-1:0] integrator_y_out;
115 logic [DATA_W-1:0] integrator_vx_out;
116 logic [DATA_W-1:0] integrator_vy_out;
117
118 function automatic logic [ACTIVE_W-1:0] ptr_to_active(input logic [PTR_W-1:0] value);
119 ptr_to_active = {{(ACTIVE_W-PTR_W){1'b0}}, value};
120 endfunction
121
122 function automatic logic [ACTIVE_W-1:0] load_idx_to_active(
123 input logic [3:0] value
124 );
125 load_idx_to_active = ACTIVE_W'(value);
126 endfunction
127
128 function automatic logic [ACTIVE_W-1:0] lane_idx_to_active(
129 input logic [PTR_W-1:0] base,
130 input logic [1:0] grp,
131 input logic [1:0] lane
132 );
133 lane_idx_to_active = ptr_to_active(base) + ACTIVE_W'({grp, 2'b00}) + ACTIVE_W'(lane)
134 ;
135 endfunction
136
137 function automatic logic is_last_active_body(
138 input logic [PTR_W-1:0] idx,
139 input logic [ACTIVE_W-1:0] active_bodies
140 );
141 is_last_active_body = (idx == MAX_BODY_PTR) ||
142 (ptr_to_active(idx) + ACTIVE_W'(1) >= active_bodies);
143 endfunction
144
145 function automatic logic [PTR_W-1:0] lane_body_idx(
146 input logic [PTR_W-1:0] base,
147 input logic [1:0] grp,
148 input logic [1:0] lane
149 );
150 lane_body_idx = base + PTR_W'({grp, 2'b00}) + PTR_W'(lane);
151 endfunction
152
153 function automatic logic lane_is_active(
154 input logic [PTR_W-1:0] base,
155 input logic [1:0] grp,
156 input logic [1:0] lane,
157 input logic [ACTIVE_W-1:0] active_bodies
158 );
159 logic [ACTIVE_W-1:0] lane_idx;
160 begin
lane_idx = lane_idx_to_active(base, grp, lane);

```

```

161     lane_is_active = (lane_idx < active_bodies);
162     end
163 endfunction
164
165 function automatic logic [3:0] make_lane_mask(
166     input logic [PTR_W-1:0] base,
167     input logic [1:0]      grp,
168     input logic [PTR_W-1:0] j_idx,
169     input logic [ACTIVE_W-1:0] active_bodies
170 );
171     logic [ACTIVE_W-1:0] lane0;
172     logic [ACTIVE_W-1:0] lane1;
173     logic [ACTIVE_W-1:0] lane2;
174     logic [ACTIVE_W-1:0] lane3;
175     logic [ACTIVE_W-1:0] j_idx_active;
176     begin
177         lane0 = lane_idx_to_active(base, grp, 2'd0);
178         lane1 = lane_idx_to_active(base, grp, 2'd1);
179         lane2 = lane_idx_to_active(base, grp, 2'd2);
180         lane3 = lane_idx_to_active(base, grp, 2'd3);
181         j_idx_active = ptr_to_active(j_idx);
182
183         make_lane_mask[0] = (lane0 >= active_bodies) || (lane0 == j_idx_active);
184         make_lane_mask[1] = (lane1 >= active_bodies) || (lane1 == j_idx_active);
185         make_lane_mask[2] = (lane2 >= active_bodies) || (lane2 == j_idx_active);
186         make_lane_mask[3] = (lane3 >= active_bodies) || (lane3 == j_idx_active);
187     end
188 endfunction
189
190 always_comb begin
191     if (n_bodies > MAX_BODIES_U32) begin
192         active_count_next = ACTIVE_W'(MAX_BODIES);
193     end else begin
194         active_count_next = n_bodies[ACTIVE_W-1:0];
195     end
196 end
197
198 always_comb begin
199     body_raddr = '0;
200
201     unique case (state)
202         ST_LOAD_TILE_PRIME: begin
203             body_raddr = tile_base + PTR_W'(load_idx_count);
204         end
205
206         ST_LOAD_TILE: begin
207             if (load_idx_count == 4'd15) begin
208                 body_raddr = tile_base + PTR_W'(load_idx_count);
209             end else begin
210                 body_raddr = tile_base + PTR_W'(load_idx_count) + PTR_W'(1);
211             end
212         end
213
214         ST_COMPUTE_PRIME: begin
215             body_raddr = j_body_idx;
216         end
217
218         ST_COMPUTE_GROUP: begin
219             if ((compute_grp == 2'd2) &&
220                 !is_last_active_body(j_body_idx, active_count)) begin
221                 body_raddr = j_body_idx + 1'b1;
222             end else begin
223                 body_raddr = j_body_idx;
224             end
225         end
226
227         ST_INTEGRATE_PRIME,
228         ST_INTEGRATE_START,
229         ST_INTEGRATE_WAIT: begin
230             body_raddr = integrate_idx;
231         end
232
233         default: begin

```

```

234         body_raddr = '0;
235     end
236 endcase
237 end
238
239 // Leapfrog starts with one half-step velocity kick. Later kicks are full-step.
240 assign integrator_half_step = run_initial_half_step && (timestep_count == 32'd0);
241 assign integrator_ax_half_exp = (body_ax[25:18] > 8'd1) ? (body_ax[25:18] - 8'd1) : 8'd0;
242 assign integrator_ay_half_exp = (body_ay[25:18] > 8'd1) ? (body_ay[25:18] - 8'd1) : 8'd0;
243 assign integrator_ax_in      = integrator_half_step ? {body_ax[26],
integrator_ax_half_exp, body_ax[17:0]} : body_ax;
244 assign integrator_ay_in      = integrator_half_step ? {body_ay[26],
integrator_ay_half_exp, body_ay[17:0]} : body_ay;
245
246 four_core_wrapper #(
247     .DATA_W(DATA_W)
248 ) u_core (
249     .i_clk      (clk),
250     .i_rst      (core_rst_n),
251     .i_clear_prev(core_clear_prev),
252     .i_load_en  (core_load_en),
253     .i_compute_en(core_compute_en),
254     .i_load_idx (core_load_idx),
255     .i_load_x   (core_load_x),
256     .i_load_y   (core_load_y),
257     .i_grp_sel  (core_grp_sel),
258     .i_j_x      (core_j_x),
259     .i_j_y      (core_j_y),
260     .i_j_m      (core_j_m),
261     .i_lane_mask(core_lane_mask),
262     .o_res0_x   (core_res_x[0]),
263     .o_res0_y   (core_res_y[0]),
264     .o_res1_x   (core_res_x[1]),
265     .o_res1_y   (core_res_y[1]),
266     .o_res2_x   (core_res_x[2]),
267     .o_res2_y   (core_res_y[2]),
268     .o_res3_x   (core_res_x[3]),
269     .o_res3_y   (core_res_y[3]),
270     .o_res_vld  (unused_core_res_vld)
271 );
272
273 nbody_integrator #(
274     .DATA_W(DATA_W)
275 ) u_integrator (
276     .clk      (clk),
277     .reset    (reset),
278     .i_start  (integrator_start),
279     .o_done   (integrator_done),
280     .i_x      (body_x),
281     .i_y      (body_y),
282     .i_vx     (body_vx),
283     .i_vy     (body_vy),
284     .i_ax     (integrator_ax_in),
285     .i_ay     (integrator_ay_in),
286     .o_x      (integrator_x_out),
287     .o_y      (integrator_y_out),
288     .o_vx     (integrator_vx_out),
289     .o_vy     (integrator_vy_out)
290 );
291
292 // Main FSM
293 always_ff @(posedge clk) begin
294     if (reset) begin
295         core_rst_n      <= 1'b0;
296         state           <= ST_IDLE;
297         update_next     <= UPD_DONE;
298         done            <= 1'b0;
299         timestep_count  <= 32'd0;
300         run_initial_half_step <= 1'b0;
301         tile_base      <= '0;
302         j_body_idx     <= '0;

```

```

303     current_j_idx      <= '0;
304     integrate_idx     <= '0;
305     wait_count        <= '0;
306     compute_grp       <= 2'd0;
307     store_lane        <= 2'd0;
308     active_count      <= '0;
309     core_clear_prev   <= 1'b0;
310     core_load_en      <= 1'b0;
311     core_compute_en   <= 1'b0;
312     load_idx_count    <= 4'd0;
313     core_load_idx     <= 4'd0;
314     core_load_x       <= '0;
315     core_load_y       <= '0;
316     core_grp_sel      <= 2'd0;
317     core_j_x          <= '0;
318     core_j_y          <= '0;
319     core_j_m          <= '0;
320     core_lane_mask    <= 4'hF;
321     current_j_x       <= '0;
322     current_j_y       <= '0;
323     current_j_m       <= '0;
324     integrator_start  <= 1'b0;
325     body_update_we    <= 1'b0;
326     body_update_addr  <= '0;
327     body_update_x     <= '0;
328     body_update_y     <= '0;
329     body_update_vx    <= '0;
330     body_update_vy    <= '0;
331     accel_we          <= 1'b0;
332     accel_waddr       <= '0;
333     accel_ax          <= '0;
334     accel_ay          <= '0;
335 end else begin
336     core_clear_prev   <= 1'b0;
337     core_load_en      <= 1'b0;
338     core_compute_en   <= 1'b0;
339     integrator_start  <= 1'b0;
340     body_update_we    <= 1'b0;
341     accel_we          <= 1'b0;
342     core_rst_n        <= 1'b1;
343
344     unique case (state)
345     ST_IDLE: begin
346         done <= 1'b0;
347         if (go) begin
348             tile_base      <= '0;
349             load_idx_count  <= 4'd0;
350             core_load_idx   <= 4'd0;
351             compute_grp     <= 2'd0;
352             j_body_idx      <= '0;
353             current_j_idx   <= '0;
354             integrate_idx   <= '0;
355             timestep_count  <= 32'd0;
356             run_initial_half_step <= first_step;
357             active_count    <= active_count_next;
358             if (active_count_next == '0) begin
359                 state <= ST_DONE;
360             end else begin
361                 state <= ST_LOAD_TILE_PRIME;
362             end
363         end
364     end
365
366     ST_LOAD_TILE_PRIME: begin
367         state <= ST_LOAD_TILE;
368     end
369
370     ST_LOAD_TILE: begin
371         core_load_en <= 1'b1;
372         core_load_idx <= load_idx_count;
373
374         if (ptr_to_active(tile_base) + load_idx_to_active(load_idx_count) <
active_count) begin

```

```

375         core_load_x <= body_x;
376         core_load_y <= body_y;
377     end else begin
378         core_load_x <= '0;
379         core_load_y <= '0;
380     end
381
382     if (load_idx_count == 4'd15) begin
383         load_idx_count <= 4'd0;
384         compute_grp <= 2'd0;
385         j_body_idx <= '0;
386         state <= ST_CLEAR_GROUP;
387     end else begin
388         load_idx_count <= load_idx_count + 1'b1;
389     end
390 end
391
392 ST_CLEAR_GROUP: begin
393     core_clear_prev <= 1'b1;
394     core_grp_sel <= 2'd0;
395     j_body_idx <= '0;
396     current_j_idx <= '0;
397     wait_count <= '0;
398     state <= ST_COMPUTE_PRIME;
399 end
400
401 ST_COMPUTE_PRIME: begin
402     current_j_x <= body_x;
403     current_j_y <= body_y;
404     current_j_m <= body_m;
405     current_j_idx <= j_body_idx;
406     compute_grp <= 2'd0;
407     state <= ST_COMPUTE_GROUP;
408 end
409
410 ST_COMPUTE_GROUP: begin
411     core_compute_en <= 1'b1;
412     core_grp_sel <= compute_grp;
413     core_j_x <= current_j_x;
414     core_j_y <= current_j_y;
415     core_j_m <= current_j_m;
416     core_lane_mask <= make_lane_mask(tile_base, compute_grp, current_j_idx,
active_count);
417
418     if (compute_grp == 2'd3) begin
419         if (is_last_active_body(j_body_idx, active_count)) begin
420             wait_count <= '0;
421             state <= ST_DRAIN_GROUP;
422         end else begin
423             j_body_idx <= j_body_idx + 1'b1;
424             current_j_idx <= j_body_idx + 1'b1;
425             current_j_x <= body_x;
426             current_j_y <= body_y;
427             current_j_m <= body_m;
428             compute_grp <= 2'd0;
429         end
430     end else begin
431         compute_grp <= compute_grp + 1'b1;
432     end
433 end
434
435 ST_DRAIN_GROUP: begin
436     core_grp_sel <= 2'd0;
437
438     if (wait_count == WAIT_W'(PIPE_LAT + 1)) begin
439         compute_grp <= 2'd0;
440         store_lane <= 2'd0;
441         state <= ST_STORE_GROUP;
442     end else begin
443         wait_count <= wait_count + 1'b1;
444     end
445 end
446

```

```

447 ST_STORE_GROUP: begin
448     core_grp_sel <= compute_grp;
449
450     if (lane_is_active(tile_base, compute_grp, store_lane, active_count))
begin
451         accel_we      <= 1'b1;
452         accel_waddr <= lane_body_idx(tile_base, compute_grp, store_lane);
453         accel_ax     <= core_res_x[store_lane];
454         accel_ay     <= core_res_y[store_lane];
455     end
456
457     if (store_lane == 2'd3) begin
458         state <= ST_NEXT_GROUP;
459     end else begin
460         store_lane <= store_lane + 1'b1;
461     end
462 end
463
464 ST_NEXT_GROUP: begin
465     if (compute_grp == 2'd3) begin
466         state <= ST_NEXT_TILE;
467     end else begin
468         compute_grp <= compute_grp + 1'b1;
469         core_grp_sel <= compute_grp + 1'b1;
470         store_lane  <= 2'd0;
471         state       <= ST_STORE_GROUP;
472     end
473 end
474
475 ST_NEXT_TILE: begin
476     if (ptr_to_active(tile_base) + ACTIVE_W'(16) >= active_count) begin
477         integrate_idx <= '0;
478         state        <= ST_INTEGRATE_PRIME;
479     end else begin
480         tile_base     <= tile_base + TILE_STRIDE;
481         load_idx_count <= 4'd0;
482         core_load_idx <= 4'd0;
483         compute_grp  <= 2'd0;
484         j_body_idx   <= '0;
485         current_j_idx <= '0;
486         state        <= ST_LOAD_TILE_PRIME;
487     end
488 end
489
490 ST_INTEGRATE_PRIME: begin
491     state <= ST_INTEGRATE_START;
492 end
493
494 ST_INTEGRATE_START: begin
495     integrator_start <= 1'b1;
496     state            <= ST_INTEGRATE_WAIT;
497 end
498
499 ST_INTEGRATE_WAIT: begin
500     if (integrator_done) begin
501         body_update_we <= 1'b1;
502         body_update_addr <= integrate_idx;
503         body_update_x   <= integrator_x_out;
504         body_update_y   <= integrator_y_out;
505         body_update_vx  <= integrator_vx_out;
506         body_update_vy  <= integrator_vy_out;
507
508         if (is_last_active_body(integrate_idx, active_count)) begin
509             if (timestep_count + 32'd1 >= gap) begin
510                 update_next <= UPD_DONE;
511             end else begin
512                 update_next <= UPD_NEXT_STEP;
513                 timestep_count <= timestep_count + 1'b1;
514             end
515         end else begin
516             update_next <= UPD_NEXT_BODY;
517             integrate_idx <= integrate_idx + 1'b1;
518         end
519     end

```

```

519         state <= ST_UPDATE_COMMIT;
520     end
521     end
522 end
523
524 ST_UPDATE_COMMIT: begin
525     unique case (update_next)
526     UPD_NEXT_BODY: begin
527         state <= ST_INTEGRATE_PRIME;
528     end
529
530     UPD_NEXT_STEP: begin
531         tile_base     <= '0;
532         load_idx_count <= 4'd0;
533         core_load_idx <= 4'd0;
534         compute_grp   <= 2'd0;
535         j_body_idx    <= '0;
536         current_j_idx <= '0;
537         integrate_idx <= '0;
538         state         <= ST_LOAD_TILE_PRIME;
539     end
540
541     default: begin
542         state <= ST_DONE;
543     end
544 endcase
545 end
546
547 ST_DONE: begin
548     done <= 1'b1;
549     if (!read_enable) begin
550         done <= 1'b0;
551         state <= ST_IDLE;
552     end
553 end
554
555 default: begin
556     state <= ST_IDLE;
557 end
558 endcase
559 end
560 end
561
562 endmodule

```

### A.3.4 four\_core\_wrapper.sv

Listing 23: code/Hardware/code/four\_core\_wrapper.sv

```

1  module four_core_wrapper #(
2      parameter int DATA_W = 27
3  ) (
4      input logic      i_clk,
5      input logic      i_rst,    // active-low synchronous reset
6
7      // external phase control
8      input logic      i_clear_prev,
9      input logic      i_load_en,
10     input logic      i_compute_en,
11
12     // load 16 i-bodies into local wrapper memory
13     input logic [3:0] i_load_idx,
14     input logic [DATA_W-1:0] i_load_x,
15     input logic [DATA_W-1:0] i_load_y,
16
17     // external group select: combinationally selects which 4 i-bodies feed the datapath
18     // or which cached 4-lane result group is read out
19     input logic [1:0] i_grp_sel,
20

```

```

21 // broadcasted j input
22 input logic [DATA_W-1:0] i_j_x,
23 input logic [DATA_W-1:0] i_j_y,
24 input logic [DATA_W-1:0] i_j_m,
25 input logic [3:0] i_lane_mask,
26
27 // accumulated outputs
28 output logic [DATA_W-1:0] o_res0_x,
29 output logic [DATA_W-1:0] o_res0_y,
30 output logic [DATA_W-1:0] o_res1_x,
31 output logic [DATA_W-1:0] o_res1_y,
32 output logic [DATA_W-1:0] o_res2_x,
33 output logic [DATA_W-1:0] o_res2_y,
34 output logic [DATA_W-1:0] o_res3_x,
35 output logic [DATA_W-1:0] o_res3_y,
36
37 // selected group results valid
38 output logic o_res_vld
39 );
40
41 localparam int PIPE_LAT = 18;
42 integer i;
43
44 // local i-body memory (16 entries)
45 logic [DATA_W-1:0] i_x_mem [16];
46 logic [DATA_W-1:0] i_y_mem [16];
47
48 // internal cached accumulated outputs (real wrapper outputs)
49 logic [DATA_W-1:0] acc_x_bank [16];
50 logic [DATA_W-1:0] acc_y_bank [16];
51 logic [3:0] grp_written;
52
53 // fixed feedback chain (no per-group selection here)
54 logic [DATA_W-1:0] prev0_x_d0, prev0_x_d1, prev0_y_d0, prev0_y_d1;
55 logic [DATA_W-1:0] prev1_x_d0, prev1_x_d1, prev1_y_d0, prev1_y_d1;
56 logic [DATA_W-1:0] prev2_x_d0, prev2_x_d1, prev2_y_d0, prev2_y_d1;
57 logic [DATA_W-1:0] prev3_x_d0, prev3_x_d1, prev3_y_d0, prev3_y_d1;
58
59 // c0->c17 pipelined delay for writeback address alignment
60 logic vld_pipe [PIPE_LAT];
61 logic [1:0] grp_pipe [PIPE_LAT];
62
63 // current group selection for datapath/readout
64 logic [3:0] grp_base;
65 logic [3:0] wb_base;
66
67 assign grp_base = {i_grp_sel, 2'b00};
68 assign wb_base = {grp_pipe[PIPE_LAT-1], 2'b00};
69
70 logic [DATA_W-1:0] cur_i0_x;
71 logic [DATA_W-1:0] cur_i0_y;
72 logic [DATA_W-1:0] cur_i1_x;
73 logic [DATA_W-1:0] cur_i1_y;
74 logic [DATA_W-1:0] cur_i2_x;
75 logic [DATA_W-1:0] cur_i2_y;
76 logic [DATA_W-1:0] cur_i3_x;
77 logic [DATA_W-1:0] cur_i3_y;
78
79 assign cur_i0_x = i_x_mem[grp_base + 4'd0];
80 assign cur_i0_y = i_y_mem[grp_base + 4'd0];
81 assign cur_i1_x = i_x_mem[grp_base + 4'd1];
82 assign cur_i1_y = i_y_mem[grp_base + 4'd1];
83 assign cur_i2_x = i_x_mem[grp_base + 4'd2];
84 assign cur_i2_y = i_y_mem[grp_base + 4'd2];
85 assign cur_i3_x = i_x_mem[grp_base + 4'd3];
86 assign cur_i3_y = i_y_mem[grp_base + 4'd3];
87
88 // active-high mask: 1 => self/invalid/hold, so new term is nulled by mj=0
89 logic [DATA_W-1:0] j0_m_eff;
90 logic [DATA_W-1:0] j1_m_eff;
91 logic [DATA_W-1:0] j2_m_eff;
92 logic [DATA_W-1:0] j3_m_eff;
93

```

```

94 assign j0_m_eff = i_lane_mask[0] ? '0 : i_j_m;
95 assign j1_m_eff = i_lane_mask[1] ? '0 : i_j_m;
96 assign j2_m_eff = i_lane_mask[2] ? '0 : i_j_m;
97 assign j3_m_eff = i_lane_mask[3] ? '0 : i_j_m;
98
99 // outputs of datapath instances
100 logic [DATA_W-1:0] dp_out0_x, dp_out0_y;
101 logic [DATA_W-1:0] dp_out1_x, dp_out1_y;
102 logic [DATA_W-1:0] dp_out2_x, dp_out2_y;
103 logic [DATA_W-1:0] dp_out3_x, dp_out3_y;
104
105 fourcore_bcj_datapath #(
106     .DATA_W(DATA_W)
107 ) u_dp (
108     .i_clk      (i_clk),
109     .i_rst      (i_rst),
110     .i_j_x      (i_compute_en ? i_j_x      : '0),
111     .i_j_y      (i_compute_en ? i_j_y      : '0),
112     .i_j0_m     (i_compute_en ? j0_m_eff   : '0),
113     .i_j1_m     (i_compute_en ? j1_m_eff   : '0),
114     .i_j2_m     (i_compute_en ? j2_m_eff   : '0),
115     .i_j3_m     (i_compute_en ? j3_m_eff   : '0),
116     .i_i0_x     (i_compute_en ? cur_i0_x   : '0),
117     .i_i0_y     (i_compute_en ? cur_i0_y   : '0),
118     .i_prev0_x (prev0_x_d1),
119     .i_prev0_y (prev0_y_d1),
120     .o_out0_x  (dp_out0_x),
121     .o_out0_y  (dp_out0_y),
122     .i_i1_x    (i_compute_en ? cur_i1_x   : '0),
123     .i_i1_y    (i_compute_en ? cur_i1_y   : '0),
124     .i_prev1_x (prev1_x_d1),
125     .i_prev1_y (prev1_y_d1),
126     .o_out1_x  (dp_out1_x),
127     .o_out1_y  (dp_out1_y),
128     .i_i2_x    (i_compute_en ? cur_i2_x   : '0),
129     .i_i2_y    (i_compute_en ? cur_i2_y   : '0),
130     .i_prev2_x (prev2_x_d1),
131     .i_prev2_y (prev2_y_d1),
132     .o_out2_x  (dp_out2_x),
133     .o_out2_y  (dp_out2_y),
134     .i_i3_x    (i_compute_en ? cur_i3_x   : '0),
135     .i_i3_y    (i_compute_en ? cur_i3_y   : '0),
136     .i_prev3_x (prev3_x_d1),
137     .i_prev3_y (prev3_y_d1),
138     .o_out3_x  (dp_out3_x),
139     .o_out3_y  (dp_out3_y)
140 );
141
142 task automatic clear_wrapper_state;
143     for (i = 0; i < 16; i = i + 1) begin
144         acc_x_bank[i] <= '0;
145         acc_y_bank[i] <= '0;
146     end
147
148     grp_written <= 4'b0000;
149
150     prev0_x_d0 <= '0; prev0_x_d1 <= '0; prev0_y_d0 <= '0; prev0_y_d1 <= '0;
151     prev1_x_d0 <= '0; prev1_x_d1 <= '0; prev1_y_d0 <= '0; prev1_y_d1 <= '0;
152     prev2_x_d0 <= '0; prev2_x_d1 <= '0; prev2_y_d0 <= '0; prev2_y_d1 <= '0;
153     prev3_x_d0 <= '0; prev3_x_d1 <= '0; prev3_y_d0 <= '0; prev3_y_d1 <= '0;
154
155     for (i = 0; i < PIPE_LAT; i = i + 1) begin
156         vld_pipe[i] <= 1'b0;
157         grp_pipe[i] <= 2'd0;
158     end
159 endtask
160
161 // sequential state updates
162 always_ff @(posedge i_clk) begin
163     if (!i_rst) begin
164         for (i = 0; i < 16; i = i + 1) begin
165             i_x_mem[i] <= '0;
166             i_y_mem[i] <= '0;

```

```

167         end
168
169         clear_wrapper_state();
170     end else if (i_load_en) begin
171         // load current i entry
172         i_x_mem[i_load_idx] <= i_load_x;
173         i_y_mem[i_load_idx] <= i_load_y;
174
175         // clear wrapper-side cached outputs / feedback / metadata
176         // during tile load so no stale values pollute the next compute phase.
177         clear_wrapper_state();
178     end else if (i_clear_prev) begin
179         // explicit clear between phases / after sufficient drain
180         clear_wrapper_state();
181     end else begin
182         // fixed feedback transport
183         prev0_x_d1 <= prev0_x_d0; prev0_y_d1 <= prev0_y_d0; prev0_x_d0 <= dp_out0_x;
184         prev0_y_d0 <= dp_out0_y;
185         prev1_x_d1 <= prev1_x_d0; prev1_y_d1 <= prev1_y_d0; prev1_x_d0 <= dp_out1_x;
186         prev1_y_d0 <= dp_out1_y;
187         prev2_x_d1 <= prev2_x_d0; prev2_y_d1 <= prev2_y_d0; prev2_x_d0 <= dp_out2_x;
188         prev2_y_d0 <= dp_out2_y;
189         prev3_x_d1 <= prev3_x_d0; prev3_y_d1 <= prev3_y_d0; prev3_x_d0 <= dp_out3_x;
190         prev3_y_d0 <= dp_out3_y;
191
192         // sample current external group as the true c0 group when this issue enters
193         core
194         vld_pipe[0] <= i_compute_en;
195         grp_pipe[0] <= i_grp_sel;
196         for (int i = 0; i < PIPE_LAT-1; i++) begin
197             vld_pipe[i+1] <= vld_pipe[i];
198             grp_pipe[i+1] <= grp_pipe[i];
199         end
200
201         // c17-aligned out_mem writeback (wrapper result storage)
202         if (vld_pipe[PIPE_LAT-1]) begin
203             acc_x_bank[wb_base + 4'd0] <= dp_out0_x;
204             acc_y_bank[wb_base + 4'd0] <= dp_out0_y;
205             acc_x_bank[wb_base + 4'd1] <= dp_out1_x;
206             acc_y_bank[wb_base + 4'd1] <= dp_out1_y;
207             acc_x_bank[wb_base + 4'd2] <= dp_out2_x;
208             acc_y_bank[wb_base + 4'd2] <= dp_out2_y;
209             acc_x_bank[wb_base + 4'd3] <= dp_out3_x;
210             acc_y_bank[wb_base + 4'd3] <= dp_out3_y;
211             grp_written[grp_pipe[PIPE_LAT-1]] <= 1'b1;
212         end
213     end
214 end
215
216 // real output is always selected from out_mem, using the same single grp_sel
217 assign o_res0_x = acc_x_bank[grp_base + 4'd0];
218 assign o_res0_y = acc_y_bank[grp_base + 4'd0];
219 assign o_res1_x = acc_x_bank[grp_base + 4'd1];
220 assign o_res1_y = acc_y_bank[grp_base + 4'd1];
221 assign o_res2_x = acc_x_bank[grp_base + 4'd2];
222 assign o_res2_y = acc_y_bank[grp_base + 4'd2];
223 assign o_res3_x = acc_x_bank[grp_base + 4'd3];
224 assign o_res3_y = acc_y_bank[grp_base + 4'd3];
225 assign o_res_vld = grp_written[i_grp_sel];
226
227 endmodule

```

### A.3.5 fourcore\_bcj\_datapath.sv

Listing 24: code/Hardware/code/fourcore\_bcj\_datapath.sv

```

1 module fourcore_bcj_datapath #(
2     parameter int DATA_W = 27
3 ) (

```

```

4   input logic      i_clk,
5   input logic      i_rst,
6
7   // broadcasted j position, per-lane j mass
8   input logic [DATA_W-1:0] i_j_x,
9   input logic [DATA_W-1:0] i_j_y,
10  input logic [DATA_W-1:0] i_j0_m,
11  input logic [DATA_W-1:0] i_j1_m,
12  input logic [DATA_W-1:0] i_j2_m,
13  input logic [DATA_W-1:0] i_j3_m,
14
15  // lane 0
16  input logic [DATA_W-1:0] i_i0_x,
17  input logic [DATA_W-1:0] i_i0_y,
18  input logic [DATA_W-1:0] i_prev0_x,
19  input logic [DATA_W-1:0] i_prev0_y,
20  output logic [DATA_W-1:0] o_out0_x,
21  output logic [DATA_W-1:0] o_out0_y,
22
23  // lane 1
24  input logic [DATA_W-1:0] i_i1_x,
25  input logic [DATA_W-1:0] i_i1_y,
26  input logic [DATA_W-1:0] i_prev1_x,
27  input logic [DATA_W-1:0] i_prev1_y,
28  output logic [DATA_W-1:0] o_out1_x,
29  output logic [DATA_W-1:0] o_out1_y,
30
31  // lane 2
32  input logic [DATA_W-1:0] i_i2_x,
33  input logic [DATA_W-1:0] i_i2_y,
34  input logic [DATA_W-1:0] i_prev2_x,
35  input logic [DATA_W-1:0] i_prev2_y,
36  output logic [DATA_W-1:0] o_out2_x,
37  output logic [DATA_W-1:0] o_out2_y,
38
39  // lane 3
40  input logic [DATA_W-1:0] i_i3_x,
41  input logic [DATA_W-1:0] i_i3_y,
42  input logic [DATA_W-1:0] i_prev3_x,
43  input logic [DATA_W-1:0] i_prev3_y,
44  output logic [DATA_W-1:0] o_out3_x,
45  output logic [DATA_W-1:0] o_out3_y
46 );
47
48 two_body_core #(
49     .DATA_W(DATA_W)
50 ) u_core0 (
51     .i_clk      (i_clk),
52     .i_rst      (i_rst),
53     .i_b1_x     (i_i0_x),
54     .i_b1_y     (i_i0_y),
55     .i_b2_x     (i_j_x),
56     .i_b2_y     (i_j_y),
57     .i_m_b2     (i_j0_m),
58     .i_a_b1_x   (i_prev0_x),
59     .i_a_b1_y   (i_prev0_y),
60     .o_a_b1_x   (o_out0_x),
61     .o_a_b1_y   (o_out0_y)
62 );
63
64 two_body_core #(
65     .DATA_W(DATA_W)
66 ) u_core1 (
67     .i_clk      (i_clk),
68     .i_rst      (i_rst),
69     .i_b1_x     (i_i1_x),
70     .i_b1_y     (i_i1_y),
71     .i_b2_x     (i_j_x),
72     .i_b2_y     (i_j_y),
73     .i_m_b2     (i_j1_m),
74     .i_a_b1_x   (i_prev1_x),
75     .i_a_b1_y   (i_prev1_y),
76     .o_a_b1_x   (o_out1_x),

```

```

77     .o_a_b1_y(o_out1_y)
78 );
79
80 two_body_core #(
81     .DATA_W(DATA_W)
82 ) u_core2 (
83     .i_clk    (i_clk),
84     .i_rst    (i_rst),
85     .i_b1_x   (i_i2_x),
86     .i_b1_y   (i_i2_y),
87     .i_b2_x   (i_j_x),
88     .i_b2_y   (i_j_y),
89     .i_m_b2   (i_j2_m),
90     .i_a_b1_x(i_prev2_x),
91     .i_a_b1_y(i_prev2_y),
92     .o_a_b1_x(o_out2_x),
93     .o_a_b1_y(o_out2_y)
94 );
95
96 two_body_core #(
97     .DATA_W(DATA_W)
98 ) u_core3 (
99     .i_clk    (i_clk),
100    .i_rst    (i_rst),
101    .i_b1_x   (i_i3_x),
102    .i_b1_y   (i_i3_y),
103    .i_b2_x   (i_j_x),
104    .i_b2_y   (i_j_y),
105    .i_m_b2   (i_j3_m),
106    .i_a_b1_x(i_prev3_x),
107    .i_a_b1_y(i_prev3_y),
108    .o_a_b1_x(o_out3_x),
109    .o_a_b1_y(o_out3_y)
110 );
111
112 endmodule

```

### A.3.6 two\_body\_core.sv

Listing 25: code/Hardware/code/two\_body\_core.sv

```

1  //=====
2  // two_body_core
3  //
4  // Timing summary:
5  //   c0      : capture b1/b2/m2 inputs
6  //   c0->c2  : dx, dy
7  //   c2->c3  : dx2, dy2
8  //   c3->c5  : r2
9  //   c5->c7  : r2 + eps^2
10 //   c7->c11 : invsqrt raw
11 //   c11->c12: register s
12 //   c12->c13: s2, m2*s
13 //   c13->c14: k = m2*s^3
14 //   c14->c15: term = dx*k, dy*k (mul comb + REG)
15 //   c15->c17: out = prev + term_reg (FpAdd = 2)
16 //
17 // prev must be driven from outside for the c15 sampling edge.
18 //=====
19
20 module two_body_core #(
21     parameter int DATA_W = 27
22 ) (
23     input logic    i_clk,
24     input logic    i_rst,    // active-low synchronous reset
25
26     // 27-bit S1E8M18 inputs
27     input logic [DATA_W-1:0] i_b1_x,
28     input logic [DATA_W-1:0] i_b1_y,

```

```

29     input logic [DATA_W-1:0] i_b2_x,
30     input logic [DATA_W-1:0] i_b2_y,
31     input logic [DATA_W-1:0] i_m_b2,
32
33     // 27-bit previous accel
34     input logic [DATA_W-1:0] i_a_b1_x,
35     input logic [DATA_W-1:0] i_a_b1_y,
36
37     // 27-bit output accel
38     output logic [DATA_W-1:0] o_a_b1_x,
39     output logic [DATA_W-1:0] o_a_b1_y
40 );
41
42     // eps^2 = 0.25
43     localparam logic [26:0] EPSILON_SQUARE = {1'b0, 8'd125, 18'd0};
44
45     // Body inputs
46     logic [26:0] b1_x_27;
47     logic [26:0] b1_y_27;
48     logic [26:0] b2_x_27;
49     logic [26:0] b2_y_27;
50     logic [26:0] m2_27;
51
52     assign b1_x_27 = i_b1_x;
53     assign b1_y_27 = i_b1_y;
54     assign b2_x_27 = i_b2_x;
55     assign b2_y_27 = i_b2_y;
56     assign m2_27   = i_m_b2;
57
58     // c0: register core-format inputs
59     logic [26:0] r_b1_x_c0;
60     logic [26:0] r_b1_y_c0;
61     logic [26:0] r_b2_x_c0;
62     logic [26:0] r_b2_y_c0;
63     logic [26:0] r_m2_c0;
64
65     always_ff @(posedge i_clk) begin
66         if (!i_rst) begin
67             r_b1_x_c0 <= 27'd0;
68             r_b1_y_c0 <= 27'd0;
69             r_b2_x_c0 <= 27'd0;
70             r_b2_y_c0 <= 27'd0;
71             r_m2_c0   <= 27'd0;
72         end else begin
73             r_b1_x_c0 <= b1_x_27;
74             r_b1_y_c0 <= b1_y_27;
75             r_b2_x_c0 <= b2_x_27;
76             r_b2_y_c0 <= b2_y_27;
77             r_m2_c0   <= m2_27;
78         end
79     end
80
81     // Reset-mux for fplib blocks (they have no reset)
82     logic en;
83     logic [26:0] b2x_in;
84     logic [26:0] b2y_in;
85
86     assign en      = i_rst;
87     assign b2x_in = en ? r_b2_x_c0 : 27'd0;
88     assign b2y_in = en ? r_b2_y_c0 : 27'd0;
89
90     // c0 comb: negate b1 so dx = x2 + (-x1), dy = y2 + (-y1)
91     logic [26:0] w_b1_x_neg;
92     logic [26:0] w_b1_y_neg;
93
94     FpNegate u_neg_b1x (
95         .iA      (en ? r_b1_x_c0 : 27'd0),
96         .oNegative(w_b1_x_neg)
97     );
98
99     FpNegate u_neg_b1y (
100        .iA      (en ? r_b1_y_c0 : 27'd0),
101        .oNegative(w_b1_y_neg)

```

```

102 );
103
104 // c0 -> c2: dx, dy (FpAdd = 2 cycles)
105 logic [26:0] w_dx_c2;
106 logic [26:0] w_dy_c2;
107
108 FpAdd u_add_dx (
109     .iCLK(i_clk),
110     .iA (b2x_in),
111     .iB (w_b1_x_neg),
112     .oSum(w_dx_c2)
113 );
114
115 FpAdd u_add_dy (
116     .iCLK(i_clk),
117     .iA (b2y_in),
118     .iB (w_b1_y_neg),
119     .oSum(w_dy_c2)
120 );
121
122 // c2 -> c3: dx2, dy2 (mul comb + reg = 1 cycle)
123 logic [26:0] w_dx2_comb;
124 logic [26:0] w_dy2_comb;
125 logic [26:0] r_dx2_c3;
126 logic [26:0] r_dy2_c3;
127
128 FpMul u_mul_dx2 (
129     .iA (w_dx_c2),
130     .iB (w_dx_c2),
131     .oProd(w_dx2_comb)
132 );
133
134 FpMul u_mul_dy2 (
135     .iA (w_dy_c2),
136     .iB (w_dy_c2),
137     .oProd(w_dy2_comb)
138 );
139
140 always_ff @(posedge i_clk) begin
141     if (!i_rst) begin
142         r_dx2_c3 <= 27'd0;
143         r_dy2_c3 <= 27'd0;
144     end else begin
145         r_dx2_c3 <= w_dx2_comb;
146         r_dy2_c3 <= w_dy2_comb;
147     end
148 end
149
150 // c3 -> c5: r2 = dx2 + dy2 (FpAdd = 2 cycles)
151 logic [26:0] w_r2_c5;
152
153 FpAdd u_add_r2 (
154     .iCLK(i_clk),
155     .iA (en ? r_dx2_c3 : 27'd0),
156     .iB (en ? r_dy2_c3 : 27'd0),
157     .oSum(w_r2_c5)
158 );
159
160 // c5 -> c7: r2e = r2 + eps2 (FpAdd = 2 cycles)
161 logic [26:0] w_r2e_c7;
162
163 FpAdd u_add_r2e (
164     .iCLK(i_clk),
165     .iA (en ? w_r2_c5 : 27'd0),
166     .iB (en ? EPSILON_SQUARE : 27'd0),
167     .oSum(w_r2e_c7)
168 );
169
170 // c7 -> c11: s_raw = inv_sqrt(r2e)
171 logic [26:0] w_s_c11;
172
173 FpInvSqrt u_invsqrt (
174     .iCLK (i_clk),

```

```

175     .iA      (en ? w_r2e_c7 : 27'd0),
176     .oInvSqrt(w_s_c11)
177 );
178
179 // c11 -> c12: register s before mul usage
180 logic [26:0] r_s_c12;
181
182 always_ff @(posedge i_clk) begin
183     if (!i_rst) begin
184         r_s_c12 <= 27'd0;
185     end else begin
186         r_s_c12 <= w_s_c11;
187     end
188 end
189
190 // Align m2 to c12
191 logic [26:0] r_m2_pipe [12];
192
193 always_ff @(posedge i_clk) begin
194     if (!i_rst) begin
195         for (int i = 0; i < 12; i++) begin
196             r_m2_pipe[i] <= 27'd0;
197         end
198     end else begin
199         r_m2_pipe[0] <= r_m2_c0;
200         for (int i = 0; i < 11; i++) begin
201             r_m2_pipe[i+1] <= r_m2_pipe[i];
202         end
203     end
204 end
205
206 logic [26:0] m2_c12;
207 assign m2_c12 = r_m2_pipe[11];
208
209 // c12 -> c13: s2 = s*s, t = m2*s
210 logic [26:0] w_s2_comb;
211 logic [26:0] w_t_comb;
212 logic [26:0] r_s2_c13;
213 logic [26:0] r_t_c13;
214
215 FpMul u_mul_s2 (
216     .iA      (r_s_c12),
217     .iB      (r_s_c12),
218     .oProd(w_s2_comb)
219 );
220
221 FpMul u_mul_t (
222     .iA      (m2_c12),
223     .iB      (r_s_c12),
224     .oProd(w_t_comb)
225 );
226
227 always_ff @(posedge i_clk) begin
228     if (!i_rst) begin
229         r_s2_c13 <= 27'd0;
230         r_t_c13 <= 27'd0;
231     end else begin
232         r_s2_c13 <= w_s2_comb;
233         r_t_c13 <= w_t_comb;
234     end
235 end
236
237 // c13 -> c14: k = m2*s^3
238 logic [26:0] w_k_comb;
239 logic [26:0] r_k_c14;
240
241 FpMul u_mul_k (
242     .iA      (r_t_c13),
243     .iB      (r_s2_c13),
244     .oProd(w_k_comb)
245 );
246
247 always_ff @(posedge i_clk) begin

```

```

248     if (!i_rst) begin
249         r_k_c14 <= 27'd0;
250     end else begin
251         r_k_c14 <= w_k_comb;
252     end
253 end
254
255 // Align dx/dy to c14
256 logic [26:0] r_dx_pipe [12];
257 logic [26:0] r_dy_pipe [12];
258
259 always_ff @(posedge i_clk) begin
260     if (!i_rst) begin
261         for (int i = 0; i < 12; i++) begin
262             r_dx_pipe[i] <= 27'd0;
263             r_dy_pipe[i] <= 27'd0;
264         end
265     end else begin
266         r_dx_pipe[0] <= w_dx_c2;
267         r_dy_pipe[0] <= w_dy_c2;
268         for (int i = 0; i < 11; i++) begin
269             r_dx_pipe[i+1] <= r_dx_pipe[i];
270             r_dy_pipe[i+1] <= r_dy_pipe[i];
271         end
272     end
273 end
274
275 logic [26:0] dx_c14;
276 logic [26:0] dy_c14;
277
278 assign dx_c14 = r_dx_pipe[11];
279 assign dy_c14 = r_dy_pipe[11];
280
281 // c14 -> c15: pair term = dx*k, dy*k (mul comb + reg)
282 logic [26:0] w_ax_term_c15;
283 logic [26:0] w_ay_term_c15;
284 logic [26:0] r_ax_term_c15;
285 logic [26:0] r_ay_term_c15;
286
287 FpMul u_mul_ax (
288     .iA    (dx_c14),
289     .iB    (r_k_c14),
290     .oProd(w_ax_term_c15)
291 );
292
293 FpMul u_mul_ay (
294     .iA    (dy_c14),
295     .iB    (r_k_c14),
296     .oProd(w_ay_term_c15)
297 );
298
299 always_ff @(posedge i_clk) begin
300     if (!i_rst) begin
301         r_ax_term_c15 <= 27'd0;
302         r_ay_term_c15 <= 27'd0;
303     end else begin
304         r_ax_term_c15 <= w_ax_term_c15;
305         r_ay_term_c15 <= w_ay_term_c15;
306     end
307 end
308
309 // c15 -> c17: out = prev + term_reg (FpAdd = 2 cycles)
310 logic [26:0] w_out_x_c17;
311 logic [26:0] w_out_y_c17;
312
313 FpAdd u_add_outx (
314     .iCLK(i_clk),
315     .iA  (en ? i_a_b1_x : 27'd0),
316     .iB  (en ? r_ax_term_c15 : 27'd0),
317     .oSum(w_out_x_c17)
318 );
319
320 FpAdd u_add_outy (

```

```

321     .iCLK(i_clk),
322     .iA (en ? i_a_b1_y : 27'd0),
323     .iB (en ? r_ay_term_c15 : 27'd0),
324     .oSum(w_out_y_c17)
325 );
326
327 assign o_a_b1_x = w_out_x_c17;
328 assign o_a_b1_y = w_out_y_c17;
329
330 endmodule

```

### A.3.7 nbody\_integrator.sv

Listing 26: code/Hardware/code/nbody\_integrator.sv

```

1 // Integration unit
2 //
3 // One-body update:
4 //   vx' = vx + ax
5 //   vy' = vy + ay
6 //   x'  = x  + vx'
7 //   y'  = y  + vy'
8 //
9 // The control unit supplies ax/ay as either a half-step or full-step
10 // acceleration increment for leapfrog integration
11
12 module nbody_integrator #(
13     parameter int DATA_W = 27
14 ) (
15     input  logic      clk,
16     input  logic      reset, // active-high reset
17
18     input  logic      i_start,
19     output logic      o_done,
20
21     input  logic [DATA_W-1:0] i_x,
22     input  logic [DATA_W-1:0] i_y,
23     input  logic [DATA_W-1:0] i_vx,
24     input  logic [DATA_W-1:0] i_vy,
25     input  logic [DATA_W-1:0] i_ax,
26     input  logic [DATA_W-1:0] i_ay,
27
28     output logic [DATA_W-1:0] o_x,
29     output logic [DATA_W-1:0] o_y,
30     output logic [DATA_W-1:0] o_vx,
31     output logic [DATA_W-1:0] o_vy
32 );
33
34 typedef enum logic [1:0] {
35     ST_IDLE,
36     ST_WAIT_V,
37     ST_WAIT_XY,
38     ST_DONE
39 } state_t;
40
41 state_t state;
42 logic [2:0] wait_count;
43
44 logic [DATA_W-1:0] x_27;
45 logic [DATA_W-1:0] y_27;
46 logic [DATA_W-1:0] vx_27;
47 logic [DATA_W-1:0] vy_27;
48 logic [DATA_W-1:0] ax_27;
49 logic [DATA_W-1:0] ay_27;
50
51 logic [DATA_W-1:0] vx_new_27;
52 logic [DATA_W-1:0] vy_new_27;
53 logic [DATA_W-1:0] x_new_27;
54 logic [DATA_W-1:0] y_new_27;

```

```

55
56   FpAdd u_add_vx (
57       .iCLK(clk),
58       .iA (vx_27),
59       .iB (ax_27),
60       .oSum(vx_new_27)
61   );
62
63   FpAdd u_add_vy (
64       .iCLK(clk),
65       .iA (vy_27),
66       .iB (ay_27),
67       .oSum(vy_new_27)
68   );
69
70   FpAdd u_add_x (
71       .iCLK(clk),
72       .iA (x_27),
73       .iB (vx_new_27),
74       .oSum(x_new_27)
75   );
76
77   FpAdd u_add_y (
78       .iCLK(clk),
79       .iA (y_27),
80       .iB (vy_new_27),
81       .oSum(y_new_27)
82   );
83
84   always_ff @(posedge clk) begin
85       if (reset) begin
86           state      <= ST_IDLE;
87           wait_count <= 3'd0;
88           o_done     <= 1'b0;
89           x_27       <= '0;
90           y_27       <= '0;
91           vx_27      <= '0;
92           vy_27      <= '0;
93           ax_27      <= '0;
94           ay_27      <= '0;
95           o_x        <= '0;
96           o_y        <= '0;
97           o_vx       <= '0;
98           o_vy       <= '0;
99       end else begin
100          o_done <= 1'b0;
101
102          unique case (state)
103              ST_IDLE: begin
104                  if (i_start) begin
105                      x_27      <= i_x;
106                      y_27      <= i_y;
107                      vx_27     <= i_vx;
108                      vy_27     <= i_vy;
109                      ax_27     <= i_ax;
110                      ay_27     <= i_ay;
111                      wait_count <= 3'd0;
112                      state     <= ST_WAIT_V;
113                  end
114              end
115
116              ST_WAIT_V: begin
117                  if (wait_count == 3'd3) begin
118                      wait_count <= 3'd0;
119                      state     <= ST_WAIT_XY;
120                  end else begin
121                      wait_count <= wait_count + 1'b1;
122                  end
123              end
124
125              ST_WAIT_XY: begin
126                  if (wait_count == 3'd3) begin
127                      o_vx <= vx_new_27;

```

```

128         o_vy <= vy_new_27;
129         o_x  <= x_new_27;
130         o_y  <= y_new_27;
131         state <= ST_DONE;
132     end else begin
133         wait_count <= wait_count + 1'b1;
134     end
135 end
136
137 ST_DONE: begin
138     o_done <= 1'b1;
139     state <= ST_IDLE;
140 end
141
142 default: begin
143     state <= ST_IDLE;
144 end
145 endcase
146 end
147 end
148
149 endmodule

```

### A.3.8 nbody\_mem.sv

Listing 27: code/Hardware/code/nbody\_mem.sv

```

1  module nbody_mem #(
2      parameter int MAX_BODIES = 1024,
3      parameter int DATA_W = 27,
4      parameter int PTR_W = $clog2(MAX_BODIES)
5  ) (
6      input  logic clk,
7
8      // CPU writes one complete body when VY_IN is committed.
9      input  logic      cpu_body_we,
10     input  logic [PTR_W-1:0] cpu_body_waddr,
11     input  logic [DATA_W-1:0] cpu_x,
12     input  logic [DATA_W-1:0] cpu_y,
13     input  logic [DATA_W-1:0] cpu_m,
14     input  logic [DATA_W-1:0] cpu_vx,
15     input  logic [DATA_W-1:0] cpu_vy,
16
17     // Shared synchronous body read port. Data is valid one clock after body_raddr is
18     // sampled.
19     input  logic [PTR_W-1:0] body_raddr,
20     output logic [DATA_W-1:0] body_x,
21     output logic [DATA_W-1:0] body_y,
22     output logic [DATA_W-1:0] body_m,
23     output logic [DATA_W-1:0] body_vx,
24     output logic [DATA_W-1:0] body_vy,
25     output logic [DATA_W-1:0] body_ax,
26     output logic [DATA_W-1:0] body_ay,
27
28     // Integrator writeback port.
29     input  logic      body_update_we,
30     input  logic [PTR_W-1:0] body_update_addr,
31     input  logic [DATA_W-1:0] body_update_x,
32     input  logic [DATA_W-1:0] body_update_y,
33     input  logic [DATA_W-1:0] body_update_vx,
34     input  logic [DATA_W-1:0] body_update_vy,
35
36     // Serialized acceleration writeback port.
37     input  logic      accel_we,
38     input  logic [PTR_W-1:0] accel_waddr,
39     input  logic [DATA_W-1:0] accel_ax,
40     input  logic [DATA_W-1:0] accel_ay
41 );

```

```

42 (* ramstyle = "M10K" *) logic [DATA_W-1:0] x_mem [0:MAX_BODIES-1];
43 (* ramstyle = "M10K" *) logic [DATA_W-1:0] y_mem [0:MAX_BODIES-1];
44 (* ramstyle = "M10K" *) logic [DATA_W-1:0] m_mem [0:MAX_BODIES-1];
45 (* ramstyle = "M10K" *) logic [DATA_W-1:0] vx_mem [0:MAX_BODIES-1];
46 (* ramstyle = "M10K" *) logic [DATA_W-1:0] vy_mem [0:MAX_BODIES-1];
47 (* ramstyle = "M10K" *) logic [DATA_W-1:0] ax_mem [0:MAX_BODIES-1];
48 (* ramstyle = "M10K" *) logic [DATA_W-1:0] ay_mem [0:MAX_BODIES-1];
49
50 logic x_we;
51 logic y_we;
52 logic m_we;
53 logic vx_we;
54 logic vy_we;
55 logic ax_we;
56 logic ay_we;
57 logic [PTR_W-1:0] x_waddr;
58 logic [PTR_W-1:0] y_waddr;
59 logic [PTR_W-1:0] m_waddr;
60 logic [PTR_W-1:0] vx_waddr;
61 logic [PTR_W-1:0] vy_waddr;
62 logic [PTR_W-1:0] ax_waddr;
63 logic [PTR_W-1:0] ay_waddr;
64 logic [DATA_W-1:0] x_wdata;
65 logic [DATA_W-1:0] y_wdata;
66 logic [DATA_W-1:0] m_wdata;
67 logic [DATA_W-1:0] vx_wdata;
68 logic [DATA_W-1:0] vy_wdata;
69 logic [DATA_W-1:0] ax_wdata;
70 logic [DATA_W-1:0] ay_wdata;
71
72 always_comb begin
73     x_we = cpu_body_we;
74     y_we = cpu_body_we;
75     m_we = cpu_body_we;
76     vx_we = cpu_body_we;
77     vy_we = cpu_body_we;
78     ax_we = cpu_body_we;
79     ay_we = cpu_body_we;
80     x_waddr = cpu_body_waddr;
81     y_waddr = cpu_body_waddr;
82     m_waddr = cpu_body_waddr;
83     vx_waddr = cpu_body_waddr;
84     vy_waddr = cpu_body_waddr;
85     ax_waddr = cpu_body_waddr;
86     ay_waddr = cpu_body_waddr;
87     x_wdata = cpu_x;
88     y_wdata = cpu_y;
89     m_wdata = cpu_m;
90     vx_wdata = cpu_vx;
91     vy_wdata = cpu_vy;
92     ax_wdata = '0;
93     ay_wdata = '0;
94
95     if (body_update_we) begin
96         x_we = 1'b1;
97         y_we = 1'b1;
98         vx_we = 1'b1;
99         vy_we = 1'b1;
100        x_waddr = body_update_addr;
101        y_waddr = body_update_addr;
102        vx_waddr = body_update_addr;
103        vy_waddr = body_update_addr;
104        x_wdata = body_update_x;
105        y_wdata = body_update_y;
106        vx_wdata = body_update_vx;
107        vy_wdata = body_update_vy;
108    end
109
110    if (accel_we) begin
111        ax_we = 1'b1;
112        ay_we = 1'b1;
113        ax_waddr = accel_waddr;
114        ay_waddr = accel_waddr;

```

```

115         ax_wdata = accel_ax;
116         ay_wdata = accel_ay;
117     end
118 end
119
120 always_ff @(posedge clk) begin
121     body_x  <= x_mem[body_raddr];
122     body_y  <= y_mem[body_raddr];
123     body_m  <= m_mem[body_raddr];
124     body_vx <= vx_mem[body_raddr];
125     body_vy <= vy_mem[body_raddr];
126     body_ax <= ax_mem[body_raddr];
127     body_ay <= ay_mem[body_raddr];
128
129     if (x_we)  x_mem[x_waddr]  <= x_wdata;
130     if (y_we)  y_mem[y_waddr]  <= y_wdata;
131     if (m_we)  m_mem[m_waddr]  <= m_wdata;
132     if (vx_we) vx_mem[vx_waddr] <= vx_wdata;
133     if (vy_we) vy_mem[vy_waddr] <= vy_wdata;
134     if (ax_we) ax_mem[ax_waddr] <= ax_wdata;
135     if (ay_we) ay_mem[ay_waddr] <= ay_wdata;
136 end
137
138 endmodule

```

### A.3.9 framebuffer\_ram.sv

Listing 28: code/Hardware/code/framebuffer\_ram.sv

```

1 // Simple dual-port 1-bit-per-pixel framebuffer storage.
2 //
3 // Port A is used by the Avalon-MM slave for CPU writes.
4 // Port B is used by the VGA scanout logic for continuous reads.
5
6 module framebuffer_ram #(
7     parameter int FB_WORDS = 9600,
8     parameter int ADDR_W   = 14
9 ) (
10     input  logic          clk,
11
12     input  logic          port_a_we,
13     input  logic [ADDR_W-1:0] port_a_addr,
14     input  logic [31:0]    port_a_writedata,
15
16     input  logic [ADDR_W-1:0] port_b_addr,
17     output logic [31:0]    port_b_readdata
18 );
19
20 (* ramstyle = "M10K" *) logic [31:0] mem [0:FB_WORDS-1];
21
22 always_ff @(posedge clk) begin
23     if (port_a_we) begin
24         mem[port_a_addr] <= port_a_writedata;
25     end
26
27     port_b_readdata <= mem[port_b_addr];
28 end
29
30 endmodule

```

### A.3.10 FpAdd.sv

Listing 29: code/Hardware/code/FpAdd.sv

```

1 /******

```

```

2  * Floating Point Adder *
3  * 2-stage pipeline *
4  *****/
5  module FpAdd (
6      input logic iCLK,
7      input logic [26:0] iA,
8      input logic [26:0] iB,
9      output logic [26:0] oSum
10 );
11
12 // Extract fields of A and B.
13 logic A_s;
14 logic [7:0] A_e;
15 logic [17:0] A_f;
16 logic B_s;
17 logic [7:0] B_e;
18 logic [17:0] B_f;
19 logic A_larger;
20
21 assign A_s = iA[26];
22 assign A_e = iA[25:18];
23 assign A_f = {1'b1, iA[17:1]};
24 assign B_s = iB[26];
25 assign B_e = iB[25:18];
26 assign B_f = {1'b1, iB[17:1]};
27
28 // Shift fractions of A and B so that they align.
29 logic [7:0] exp_diff_A;
30 logic [7:0] exp_diff_B;
31 logic [7:0] larger_exp;
32 logic [36:0] A_f_shifted;
33 logic [36:0] B_f_shifted;
34
35 assign exp_diff_A = B_e - A_e; // if B bigger
36 assign exp_diff_B = A_e - B_e; // if A bigger
37
38 assign larger_exp = (B_e > A_e) ? B_e : A_e;
39
40 assign A_f_shifted = A_larger ? {1'b0, A_f, 18'b0} :
41 (exp_diff_A > 8'd35) ? 37'b0 :
42 ({1'b0, A_f, 18'b0} >> exp_diff_A);
43 assign B_f_shifted = ~A_larger ? {1'b0, B_f, 18'b0} :
44 (exp_diff_B > 8'd35) ? 37'b0 :
45 ({1'b0, B_f, 18'b0} >> exp_diff_B);
46
47 // Determine which of A, B is larger.
48 assign A_larger = (A_e > B_e) ? 1'b1 :
49 ((A_e == B_e) && (A_f > B_f)) ? 1'b1 :
50 1'b0;
51
52 // Calculate sum or difference of shifted fractions.
53 logic [36:0] pre_sum;
54 assign pre_sum = ((A_s ^ B_s) & A_larger) ? A_f_shifted - B_f_shifted :
55 ((A_s ^ B_s) & ~A_larger) ? B_f_shifted - A_f_shifted :
56 A_f_shifted + B_f_shifted;
57
58 // Buffer midway results.
59 logic [36:0] buf_pre_sum;
60 logic [7:0] buf_larger_exp;
61 logic buf_A_e_zero;
62 logic buf_B_e_zero;
63 logic [26:0] buf_A;
64 logic [26:0] buf_B;
65 logic buf_oSum_s;
66
67 always_ff @(posedge iCLK) begin
68     buf_pre_sum <= pre_sum;
69     buf_larger_exp <= larger_exp;
70     buf_A_e_zero <= (A_e == 8'b0);
71     buf_B_e_zero <= (B_e == 8'b0);
72     buf_A <= iA;
73     buf_B <= iB;
74     buf_oSum_s <= A_larger ? A_s : B_s;

```

```

75     end
76
77     // Convert to positive fraction and a sign bit.
78     logic [36:0] pre_frac;
79     assign pre_frac = buf_pre_sum;
80
81     // Determine output fraction and exponent change with position of first 1.
82     logic [17:0] oSum_f;
83     logic [7:0]  shft_amt;
84
85     assign shft_amt = pre_frac[36] ? 8'd0  : pre_frac[35] ? 8'd1  :
86                    pre_frac[34] ? 8'd2  : pre_frac[33] ? 8'd3  :
87                    pre_frac[32] ? 8'd4  : pre_frac[31] ? 8'd5  :
88                    pre_frac[30] ? 8'd6  : pre_frac[29] ? 8'd7  :
89                    pre_frac[28] ? 8'd8  : pre_frac[27] ? 8'd9  :
90                    pre_frac[26] ? 8'd10 : pre_frac[25] ? 8'd11 :
91                    pre_frac[24] ? 8'd12 : pre_frac[23] ? 8'd13 :
92                    pre_frac[22] ? 8'd14 : pre_frac[21] ? 8'd15 :
93                    pre_frac[20] ? 8'd16 : pre_frac[19] ? 8'd17 :
94                    pre_frac[18] ? 8'd18 : pre_frac[17] ? 8'd19 :
95                    pre_frac[16] ? 8'd20 : pre_frac[15] ? 8'd21 :
96                    pre_frac[14] ? 8'd22 : pre_frac[13] ? 8'd23 :
97                    pre_frac[12] ? 8'd24 : pre_frac[11] ? 8'd25 :
98                    pre_frac[10] ? 8'd26 : pre_frac[9]  ? 8'd27 :
99                    pre_frac[8]  ? 8'd28 : pre_frac[7]  ? 8'd29 :
100                   pre_frac[6]  ? 8'd30 : pre_frac[5]  ? 8'd31 :
101                   pre_frac[4]  ? 8'd32 : pre_frac[3]  ? 8'd33 :
102                   pre_frac[2]  ? 8'd34 : pre_frac[1]  ? 8'd35 :
103                   pre_frac[0]  ? 8'd36 : 8'd37;
104
105     logic [53:0] pre_frac_shft;
106     logic [53:0] uflow_shift;
107
108     // The shift +1 is because high order bit is not stored, but implied.
109     assign pre_frac_shft = {pre_frac, 17'b0} << (shft_amt + 8'd1);
110     assign uflow_shift   = {pre_frac, 17'b0} << shft_amt;
111     assign oSum_f        = pre_frac_shft[53:36];
112
113     logic [7:0] oSum_e;
114     assign oSum_e = buf_larger_exp - shft_amt + 8'd1;
115
116     // Detect underflow. If top bit of mantissa is not set, then denormalize
117     logic underflow;
118     assign underflow = ~uflow_shift[53];
119
120     always_ff @(posedge iCLK) begin
121         oSum <= (buf_A_e_zero && buf_B_e_zero) ? 27'b0 :
122                buf_A_e_zero      ? buf_B :
123                buf_B_e_zero      ? buf_A :
124                underflow          ? 27'b0 :
125                (pre_frac == 37'b0) ? 27'b0 :
126                {buf_oSum_s, oSum_e, oSum_f};
127     end
128
129 endmodule

```

### A.3.11 FpMul.sv

Listing 30: code/Hardware/code/FpMul.sv

```

1  /*****
2  * Floating Point Multiplier
3  * Fully Combinational
4  *****/
5  module FpMul (
6      input  logic [26:0] iA,    // First input
7      input  logic [26:0] iB,    // Second input
8      output logic [26:0] oProd // Product
9  );

```

```

10
11 // Extract fields of A and B.
12 logic          A_s;
13 logic [7:0]    A_e;
14 logic [17:0]   A_f;
15 logic          B_s;
16 logic [7:0]    B_e;
17 logic [17:0]   B_f;
18
19 assign A_s = iA[26];
20 assign A_e = iA[25:18];
21 assign A_f = {1'b1, iA[17:1]};
22 assign B_s = iB[26];
23 assign B_e = iB[25:18];
24 assign B_f = {1'b1, iB[17:1]};
25
26 // XOR sign bits to determine product sign.
27 logic oProd_s;
28 assign oProd_s = A_s ^ B_s;
29
30 // Multiply the fractions of A and B.
31 logic [35:0] pre_prod_frac;
32 assign pre_prod_frac = A_f * B_f;
33
34 // Add exponents of A and B.
35 logic [8:0] pre_prod_exp;
36 assign pre_prod_exp = A_e + B_e;
37
38 // If top bit of product frac is 0, shift left one.
39 logic [7:0] oProd_e;
40 logic [17:0] oProd_f;
41
42 assign oProd_e = pre_prod_frac[35] ? (pre_prod_exp - 9'd126) :
43                                     (pre_prod_exp - 9'd127);
44 assign oProd_f = pre_prod_frac[35] ? pre_prod_frac[34:17] :
45                                     pre_prod_frac[33:16];
46
47 // Detect underflow.
48 logic underflow;
49 assign underflow = pre_prod_exp < 9'h80;
50
51 // Detect zero conditions (either product frac doesn't start with 1, or underflow).
52 assign oProd = underflow      ? 27'b0 :
53                 (B_e == 8'd0) ? 27'b0 :
54                 (A_e == 8'd0) ? 27'b0 :
55                 {oProd_s, oProd_e, oProd_f};
56
57 endmodule

```

### A.3.12 FpInvSqrt.sv

Listing 31: code/Hardware/code/FpInvSqrt.sv

```

1  /*****
2  * Floating Point Fast Inverse Square Root          *
3  * 5-stage pipeline                                *
4  * http://en.wikipedia.org/wiki/Fast\_inverse\_square\_root *
5  * Magic number 27'd49920718                       *
6  * 1.5 = 27'd33423360                               *
7  *****/
8  module FpInvSqrt (
9      input  logic    iCLK,
10     input  logic [26:0] iA,
11     output logic [26:0] oInvSqrt
12 );
13
14 // Extract fields of A.
15 logic          A_s;
16 logic [7:0]    A_e;

```

```

17 logic [17:0] A_f;
18
19 assign A_s = iA[26];
20 assign A_e = iA[25:18];
21 assign A_f = iA[17:0];
22
23 // Stage 1.
24 logic [26:0] y_1;
25 logic [26:0] y_1_out;
26 logic [26:0] half_iA_1;
27
28 assign y_1 = 27'd49920718 - (iA >> 1);
29 assign half_iA_1 = {A_s, A_e - 8'd1, A_f};
30
31 FpMul s1_mult (
32     .iA (y_1),
33     .iB (y_1),
34     .oProd(y_1_out)
35 );
36
37 // Stage 2.
38 logic [26:0] y_2;
39 logic [26:0] mult_2_in;
40 logic [26:0] half_iA_2;
41 logic [26:0] y_2_out;
42
43 FpMul s2_mult (
44     .iA (half_iA_2),
45     .iB (mult_2_in),
46     .oProd(y_2_out)
47 );
48
49 // Stage 3.
50 logic [26:0] y_3;
51 logic [26:0] add_3_in;
52 logic [26:0] y_3_out;
53
54 FpAdd s3_add (
55     .iCLK(iCLK),
56     .iA ({~add_3_in[26], add_3_in[25:0]}),
57     .iB (27'd33423360),
58     .oSum(y_3_out)
59 );
60
61 // Stage 4.
62 logic [26:0] y_4;
63
64 // Stage 5.
65 logic [26:0] y_5;
66
67 FpMul s5_mult (
68     .iA (y_5),
69     .iB (y_3_out),
70     .oProd(oInvSqrt)
71 );
72
73 always_ff @(posedge iCLK) begin
74     // Stage 1 to 2.
75     y_2 <= y_1;
76     mult_2_in <= y_1_out;
77     half_iA_2 <= half_iA_1;
78
79     // Stage 2 to 3.
80     y_3 <= y_2;
81     add_3_in <= y_2_out;
82
83     // Stage 3 to 4.
84     y_4 <= y_3;
85
86     // Stage 4 to 5.
87     y_5 <= y_4;
88 end
89

```

```
90 endmodule
```

### A.3.13 FpNegate.sv

Listing 32: code/Hardware/code/FpNegate.sv

```
1  /*****
2  * Floating Point sign negation
3  * Combinational
4  *****/
5  module FpNegate (
6      input logic [26:0] iA,
7      output logic [26:0] oNegative
8  );
9
10     // Extract fields of A.
11     logic A_s;
12     logic [7:0] A_e;
13     logic [17:0] A_f;
14
15     assign A_s = iA[26];
16     assign A_e = iA[25:18];
17     assign A_f = iA[17:0];
18
19     // Flip bit 26.
20     assign oNegative = {~A_s, A_e, A_f};
21
22 endmodule
```

## A.4 Hardware Testbenches

### A.4.1 tb\_nbody\_control.sv

Listing 33: code/Hardware/tb/tb\_nbody\_control.sv

```
1  `timescale 1ns/1ps
2
3  module tb_nbody_control;
4
5      localparam int MAX_BODIES = 1024;
6      localparam int DATA_W = 27;
7      localparam int PTR_W = $clog2(MAX_BODIES);
8      localparam int RUN_TIMEOUT_CYCLES = 400000;
9
10     localparam string INPUT_FILE = "tb/frame_input/frame0_1024binit200_27bits.txt";
11     localparam string OUT_FILE = "tb/output/control_accel_1024binit200_27bits.txt";
12
13     logic clk;
14     logic reset;
15
16     logic go;
17     logic read_enable;
18     logic first_step;
19     logic [31:0] n_bodies;
20     logic [31:0] gap;
21     logic done;
22
23     logic [PTR_W-1:0] body_raddr;
24     logic [DATA_W-1:0] body_x;
25     logic [DATA_W-1:0] body_y;
26     logic [DATA_W-1:0] body_m;
27     logic [DATA_W-1:0] body_vx;
28     logic [DATA_W-1:0] body_vy;
29     logic [DATA_W-1:0] body_ax;
```

```

30 logic [DATA_W-1:0] body_ay;
31
32 logic body_update_we;
33 logic [PTR_W-1:0] body_update_addr;
34 logic [DATA_W-1:0] body_update_x;
35 logic [DATA_W-1:0] body_update_y;
36 logic [DATA_W-1:0] body_update_vx;
37 logic [DATA_W-1:0] body_update_vy;
38
39 logic accel_we;
40 logic [PTR_W-1:0] accel_waddr;
41 logic [DATA_W-1:0] accel_ax;
42 logic [DATA_W-1:0] accel_ay;
43
44 logic cpu_body_we;
45 logic [PTR_W-1:0] cpu_body_waddr;
46 logic [DATA_W-1:0] cpu_x;
47 logic [DATA_W-1:0] cpu_y;
48 logic [DATA_W-1:0] cpu_m;
49 logic [DATA_W-1:0] cpu_vx;
50 logic [DATA_W-1:0] cpu_vy;
51
52 logic [DATA_W-1:0] init_x [0:MAX_BODIES-1];
53 logic [DATA_W-1:0] init_y [0:MAX_BODIES-1];
54 logic [DATA_W-1:0] init_m [0:MAX_BODIES-1];
55 logic [DATA_W-1:0] init_vx [0:MAX_BODIES-1];
56 logic [DATA_W-1:0] init_vy [0:MAX_BODIES-1];
57 logic [DATA_W-1:0] final_x [0:MAX_BODIES-1];
58 logic [DATA_W-1:0] final_y [0:MAX_BODIES-1];
59 logic [DATA_W-1:0] final_vx [0:MAX_BODIES-1];
60 logic [DATA_W-1:0] final_vy [0:MAX_BODIES-1];
61 logic [DATA_W-1:0] final_ax [0:MAX_BODIES-1];
62 logic [DATA_W-1:0] final_ay [0:MAX_BODIES-1];
63 logic accel_seen [0:MAX_BODIES-1];
64 logic update_seen [0:MAX_BODIES-1];
65
66 int err_count;
67 int accel_count;
68 int update_count;
69 int fo;
70
71 initial clk = 1'b0;
72 always #5 clk = ~clk;
73
74 nbody_control #(
75     .MAX_BODIES(MAX_BODIES),
76     .DATA_W(DATA_W)
77 ) dut_control (
78     .clk (clk),
79     .reset (reset),
80     .go (go),
81     .read_enable (read_enable),
82     .first_step (first_step),
83     .n_bodies (n_bodies),
84     .gap (gap),
85     .done (done),
86     .body_raddr (body_raddr),
87     .body_x (body_x),
88     .body_y (body_y),
89     .body_m (body_m),
90     .body_vx (body_vx),
91     .body_vy (body_vy),
92     .body_ax (body_ax),
93     .body_ay (body_ay),
94     .body_update_we (body_update_we),
95     .body_update_addr (body_update_addr),
96     .body_update_x (body_update_x),
97     .body_update_y (body_update_y),
98     .body_update_vx (body_update_vx),
99     .body_update_vy (body_update_vy),
100    .accel_we (accel_we),
101    .accel_waddr (accel_waddr),
102    .accel_ax (accel_ax),

```

```

103     .accel_ay      (accel_ay)
104 );
105
106 nbody_mem #(
107     .MAX_BODIES(MAX_BODIES),
108     .DATA_W(DATA_W),
109     .PTR_W(PTR_W)
110 ) dut_mem (
111     .clk            (clk),
112     .cpu_body_we   (cpu_body_we),
113     .cpu_body_waddr (cpu_body_waddr),
114     .cpu_x         (cpu_x),
115     .cpu_y         (cpu_y),
116     .cpu_m         (cpu_m),
117     .cpu_vx        (cpu_vx),
118     .cpu_vy        (cpu_vy),
119     .body_raddr    (body_raddr),
120     .body_x        (body_x),
121     .body_y        (body_y),
122     .body_m        (body_m),
123     .body_vx       (body_vx),
124     .body_vy       (body_vy),
125     .body_ax       (body_ax),
126     .body_ay       (body_ay),
127     .body_update_we (body_update_we),
128     .body_update_addr (body_update_addr),
129     .body_update_x  (body_update_x),
130     .body_update_y  (body_update_y),
131     .body_update_vx (body_update_vx),
132     .body_update_vy (body_update_vy),
133     .accel_we       (accel_we),
134     .accel_waddr    (accel_waddr),
135     .accel_ax       (accel_ax),
136     .accel_ay       (accel_ay)
137 );
138
139 task automatic clear_inputs;
140     begin
141         go = 1'b0;
142         read_enable = 1'b1;
143         first_step = 1'b1;
144         n_bodies = MAX_BODIES;
145         gap = 32'd1;
146         cpu_body_we = 1'b0;
147         cpu_body_waddr = '0;
148         cpu_x = '0;
149         cpu_y = '0;
150         cpu_m = '0;
151         cpu_vx = '0;
152         cpu_vy = '0;
153     end
154 endtask
155
156 task automatic read_frame_file(input string fname);
157     int fd;
158     string line;
159     int idx;
160     logic [DATA_W-1:0] lpx, lpy, lvx, lvy, lm;
161     int got;
162     begin
163         for (int i = 0; i < MAX_BODIES; i++) begin
164             init_x[i] = '0;
165             init_y[i] = '0;
166             init_m[i] = '0;
167             init_vx[i] = '0;
168             init_vy[i] = '0;
169             final_x[i] = '0;
170             final_y[i] = '0;
171             final_vx[i] = '0;
172             final_vy[i] = '0;
173             final_ax[i] = '0;
174             final_ay[i] = '0;
175             accel_seen[i] = 1'b0;

```

```

176     update_seen[i] = 1'b0;
177     end
178
179     fd = $fopen(fname, "r");
180     if (fd == 0) $fatal(1, "ERROR: cannot open INPUT_FILE=%s", fname);
181
182     while (!$feof(fd)) begin
183         line = "";
184         void'($fgets(line, fd));
185
186         if (line.len() == 0) continue;
187         if (line.substr(0, 0) == "#") continue;
188
189         got = $sscanf(line, "%d %h %h %h %h %h",
190                     idx, lpx, lpy, lvx, lvy, lm);
191
192         if (got == 6 && idx >= 0 && idx < MAX_BODIES) begin
193             init_x[idx] = lpx;
194             init_y[idx] = lpy;
195             init_vx[idx] = lvx;
196             init_vy[idx] = lvy;
197             init_m[idx] = lm;
198         end
199     end
200
201     $fclose(fd);
202     end
203 endtask
204
205 task automatic cpu_write_body(
206     input [PTR_W-1:0] addr,
207     input [DATA_W-1:0] x,
208     input [DATA_W-1:0] y,
209     input [DATA_W-1:0] mass,
210     input [DATA_W-1:0] vx,
211     input [DATA_W-1:0] vy
212 );
213     begin
214         @(negedge clk);
215         cpu_body_we = 1'b1;
216         cpu_body_waddr = addr;
217         cpu_x = x;
218         cpu_y = y;
219         cpu_m = mass;
220         cpu_vx = vx;
221         cpu_vy = vy;
222         @(posedge clk);
223         @(negedge clk);
224         cpu_body_we = 1'b0;
225     end
226 endtask
227
228 task automatic preload_memory;
229     begin
230         for (int i = 0; i < MAX_BODIES; i++) begin
231             cpu_write_body(PTR_W'(i), init_x[i], init_y[i], init_m[i], init_vx[i], init_vy[i]);
232         end
233     end
234 endtask
235
236 task automatic pulse_go;
237     begin
238         @(negedge clk);
239         go = 1'b1;
240         @(posedge clk);
241         @(negedge clk);
242         go = 1'b0;
243     end
244 endtask
245
246 task automatic run_and_capture;
247     int cycles;
248     begin

```

```

249     cycles = 0;
250     accel_count = 0;
251     update_count = 0;
252
253     pulse_go();
254
255     while (done != 1'b1 && cycles < RUN_TIMEOUT_CYCLES) begin
256         @(posedge clk);
257         #1;
258         cycles++;
259
260         if (accel_we) begin
261             if (accel_waddr >= MAX_BODIES) begin
262                 $display("ERROR accel_waddr out of range: %0d", accel_waddr);
263                 err_count++;
264             end else begin
265                 if (accel_seen[accel_waddr]) begin
266                     $display("ERROR duplicate accel write addr=%0d", accel_waddr);
267                     err_count++;
268                 end
269                 accel_seen[accel_waddr] = 1'b1;
270                 final_ax[accel_waddr] = accel_ax;
271                 final_ay[accel_waddr] = accel_ay;
272                 accel_count++;
273             end
274         end
275
276         if (body_update_we) begin
277             if (body_update_addr >= MAX_BODIES) begin
278                 $display("ERROR body_update_addr out of range: %0d", body_update_addr);
279                 err_count++;
280             end else begin
281                 if (update_seen[body_update_addr]) begin
282                     $display("ERROR duplicate body update addr=%0d", body_update_addr);
283                     err_count++;
284                 end
285                 update_seen[body_update_addr] = 1'b1;
286                 final_x[body_update_addr] = body_update_x;
287                 final_y[body_update_addr] = body_update_y;
288                 final_vx[body_update_addr] = body_update_vx;
289                 final_vy[body_update_addr] = body_update_vy;
290                 update_count++;
291             end
292         end
293     end
294
295     if (done != 1'b1) begin
296         $display("ERROR timed out waiting for done after %0d cycles", cycles);
297         err_count++;
298     end
299
300     if (accel_count != MAX_BODIES) begin
301         $display("ERROR accel_count=%0d expected %0d", accel_count, MAX_BODIES);
302         err_count++;
303     end
304
305     if (update_count != MAX_BODIES) begin
306         $display("ERROR update_count=%0d expected %0d", update_count, MAX_BODIES);
307         err_count++;
308     end
309
310     for (int i = 0; i < MAX_BODIES; i++) begin
311         if (!accel_seen[i]) begin
312             if (err_count < 32) begin
313                 $display("ERROR missing accel write addr=%0d", i);
314             end
315             err_count++;
316         end
317
318         if (!update_seen[i]) begin
319             if (err_count < 32) begin
320                 $display("ERROR missing body update addr=%0d", i);
321             end

```

```

322     err_count++;
323     end
324 end
325
326 fo = $fopen(OUT_FILE, "w");
327 if (fo == 0) $fatal(1, "ERROR: cannot open OUT_FILE=%s", OUT_FILE);
328 $fwrite(fo, "# i x y vx vy ax ay (S1E8M18 hex)\n");
329
330 for (int i = 0; i < MAX_BODIES; i++) begin
331     $fwrite(fo, "%4d %07h %07h %07h %07h %07h %07h\n",
332           i, final_x[i], final_y[i], final_vx[i], final_vy[i],
333           final_ax[i], final_ay[i]);
334 end
335
336 $fclose(fo);
337
338 $display("DONE. cycles=%0d accel_count=%0d update_count=%0d wrote %s",
339         cycles, accel_count, update_count, OUT_FILE);
340 end
341 endtask
342
343 initial begin
344     err_count = 0;
345     reset = 1'b1;
346     clear_inputs();
347     read_frame_file(INPUT_FILE);
348
349     repeat (4) @(posedge clk);
350     @(negedge clk);
351     reset = 1'b0;
352
353     repeat (2) @(posedge clk);
354     preload_memory();
355     repeat (2) @(posedge clk);
356
357     run_and_capture();
358
359     if (err_count == 0) begin
360         $display("PASS: tb_nbody_control completed without errors");
361     end else begin
362         $display("FAIL: tb_nbody_control saw %0d error(s)", err_count);
363     end
364
365     $finish;
366 end
367
368 endmodule

```

#### A.4.2 tb\_four\_core\_wrapper.sv

Listing 34: code/Hardware/tb/tb\_four\_core\_wrapper.sv

```

1  `timescale 1ns/1ps
2
3  // Full-frame testbench for four_core_wrapper.
4  //
5  // This drives the wrapper the same way nbody_control intends to use it:
6  //   - load one 16-body tile into wrapper local memory
7  //   - clear wrapper accumulation state once for the tile
8  //   - for each j body, hold j stable while groups 0,1,2,3 are issued
9  //   - after the pipeline drains, sample all 16 accumulated lane outputs
10 //
11 // Output format:
12 //   # i ax ay (S1E8M18 hex)
13
14 module tb_four_core_wrapper;
15
16     localparam int DATA_W = 27;
17     localparam int N_BODIES = 1024;

```

```

18 localparam int TILE_SIZE = 16;
19 localparam int GROUP_SIZE = 4;
20 localparam int N_GROUPS = TILE_SIZE / GROUP_SIZE;
21 localparam int PIPE_LAT = 18;
22 localparam int RUN_TIMEOUT_CYCLES = 2000000;
23
24 localparam string INPUT_FILE = "tb/frame_input/frame0_1024binit200_27bits.txt";
25 localparam string OUT_FILE = "tb/output/four_core_wrapper_accel_1024binit200_27bits.txt"
26 ;
27
28 logic clk;
29 logic rst_n;
30 logic i_clear_prev;
31 logic i_load_en;
32 logic i_compute_en;
33 logic [3:0] i_load_idx;
34 logic [DATA_W-1:0] i_load_x;
35 logic [DATA_W-1:0] i_load_y;
36 logic [1:0] i_grp_sel;
37 logic [DATA_W-1:0] i_j_x;
38 logic [DATA_W-1:0] i_j_y;
39 logic [DATA_W-1:0] i_j_m;
40 logic [3:0] i_lane_mask;
41
42 wire [DATA_W-1:0] o_res0_x;
43 wire [DATA_W-1:0] o_res0_y;
44 wire [DATA_W-1:0] o_res1_x;
45 wire [DATA_W-1:0] o_res1_y;
46 wire [DATA_W-1:0] o_res2_x;
47 wire [DATA_W-1:0] o_res2_y;
48 wire [DATA_W-1:0] o_res3_x;
49 wire [DATA_W-1:0] o_res3_y;
50 wire o_res_vld;
51
52 logic [DATA_W-1:0] px [0:N_BODIES-1];
53 logic [DATA_W-1:0] py [0:N_BODIES-1];
54 logic [DATA_W-1:0] m [0:N_BODIES-1];
55 logic [DATA_W-1:0] ax [0:N_BODIES-1];
56 logic [DATA_W-1:0] ay [0:N_BODIES-1];
57
58 int err_count;
59 int fo;
60 int cycles;
61
62 initial clk = 1'b0;
63 always #5 clk = ~clk;
64
65 four_core_wrapper #(
66     .DATA_W(DATA_W)
67 ) dut (
68     .i_clk (clk),
69     .i_rst (rst_n),
70     .i_clear_prev(i_clear_prev),
71     .i_load_en (i_load_en),
72     .i_compute_en(i_compute_en),
73     .i_load_idx (i_load_idx),
74     .i_load_x (i_load_x),
75     .i_load_y (i_load_y),
76     .i_grp_sel (i_grp_sel),
77     .i_j_x (i_j_x),
78     .i_j_y (i_j_y),
79     .i_j_m (i_j_m),
80     .i_lane_mask (i_lane_mask),
81     .o_res0_x (o_res0_x),
82     .o_res0_y (o_res0_y),
83     .o_res1_x (o_res1_x),
84     .o_res1_y (o_res1_y),
85     .o_res2_x (o_res2_x),
86     .o_res2_y (o_res2_y),
87     .o_res3_x (o_res3_x),
88     .o_res3_y (o_res3_y),
89     .o_res_vld (o_res_vld)
90 );

```

```

90
91 task automatic drive_idle;
92     begin
93         i_clear_prev = 1'b0;
94         i_load_en = 1'b0;
95         i_compute_en = 1'b0;
96         i_load_idx = 4'd0;
97         i_load_x = '0;
98         i_load_y = '0;
99         i_j_x = '0;
100        i_j_y = '0;
101        i_j_m = '0;
102        i_lane_mask = 4'h0;
103    end
104 endtask
105
106 task automatic read_frame_file(input string fname);
107     int fd;
108     string line;
109     int idx;
110     logic [DATA_W-1:0] lpx, lpy, lvx, lvy, lm;
111     int got;
112     begin
113         for (int t = 0; t < N_BODIES; t++) begin
114             px[t] = '0;
115             py[t] = '0;
116             m[t] = '0;
117             ax[t] = '0;
118             ay[t] = '0;
119         end
120
121         fd = $fopen(fname, "r");
122         if (fd == 0) $fatal(1, "ERROR: cannot open INPUT_FILE=%s", fname);
123
124         while (!$feof(fd)) begin
125             line = "";
126             void'($fgets(line, fd));
127
128             if (line.len() == 0) continue;
129             if (line.substr(0, 0) == "#") continue;
130
131             got = $sscanf(line, "%d %h %h %h %h %h",
132                          idx, lpx, lpy, lvx, lvy, lm);
133
134             if (got == 6 && idx >= 0 && idx < N_BODIES) begin
135                 px[idx] = lpx;
136                 py[idx] = lpy;
137                 m[idx] = lm;
138             end
139         end
140
141         $fclose(fd);
142     end
143 endtask
144
145 task automatic load_tile(input int tile_base);
146     int lane;
147     int body_idx;
148     begin
149         for (lane = 0; lane < TILE_SIZE; lane++) begin
150             body_idx = tile_base + lane;
151             @(negedge clk);
152             drive_idle();
153             i_load_en = 1'b1;
154             i_load_idx = lane[3:0];
155             if (body_idx < N_BODIES) begin
156                 i_load_x = px[body_idx];
157                 i_load_y = py[body_idx];
158             end else begin
159                 i_load_x = '0;
160                 i_load_y = '0;
161             end
162             @(posedge clk);

```

```

163     end
164
165     @(negedge clk);
166     drive_idle();
167     end
168 endtask
169
170 task automatic pulse_clear_prev;
171     begin
172     @(negedge clk);
173     drive_idle();
174     i_grp_sel = 2'd0;
175     i_clear_prev = 1'b1;
176     @(posedge clk);
177     @(negedge clk);
178     i_clear_prev = 1'b0;
179     end
180 endtask
181
182 function automatic logic [3:0] make_lane_mask(input int tile_base, input int grp, input
183 int j_idx);
184     int lane_body;
185     begin
186     for (int lane = 0; lane < GROUP_SIZE; lane++) begin
187         lane_body = tile_base + grp * GROUP_SIZE + lane;
188         make_lane_mask[lane] = (lane_body >= N_BODIES) || (lane_body == j_idx);
189     end
190 endfunction
191
192 task automatic issue_one_compute(
193     input [1:0] grp_sel,
194     input [3:0] lane_mask,
195     input [DATA_W-1:0] jx,
196     input [DATA_W-1:0] jy,
197     input [DATA_W-1:0] jm
198 );
199     begin
200     @(negedge clk);
201     drive_idle();
202     i_grp_sel = grp_sel;
203     i_lane_mask = lane_mask;
204     i_j_x = jx;
205     i_j_y = jy;
206     i_j_m = jm;
207     i_compute_en = 1'b1;
208     @(posedge clk);
209     end
210 endtask
211
212 task automatic drain_tile;
213     begin
214     @(negedge clk);
215     drive_idle();
216     i_grp_sel = 2'd0;
217
218     for (int cyc = 0; cyc < PIPE_LAT + 4; cyc++) begin
219         @(posedge clk);
220     end
221
222     for (int grp = 0; grp < N_GROUPS; grp++) begin
223         @(negedge clk);
224         drive_idle();
225         i_grp_sel = grp[1:0];
226         #1;
227         if (o_res_vld != 1'b1) begin
228             $display("ERROR: o_res_vld not set for tile group %0d at t=%0t", grp, $time);
229             err_count++;
230         end
231     end
232     end
233 endtask
234

```

```

235 task automatic save_group(input int tile_base, input int grp);
236     int base;
237     begin
238         base = tile_base + grp * GROUP_SIZE;
239         if (base + 0 < N_BODIES) begin ax[base + 0] = o_res0_x; ay[base + 0] = o_res0_y; end
240         if (base + 1 < N_BODIES) begin ax[base + 1] = o_res1_x; ay[base + 1] = o_res1_y; end
241         if (base + 2 < N_BODIES) begin ax[base + 2] = o_res2_x; ay[base + 2] = o_res2_y; end
242         if (base + 3 < N_BODIES) begin ax[base + 3] = o_res3_x; ay[base + 3] = o_res3_y; end
243     end
244 endtask
245
246 task automatic run_tile(input int tile_base);
247     begin
248         load_tile(tile_base);
249         repeat (2) @(posedge clk);
250         pulse_clear_prev();
251
252         for (int j = 0; j < N_BODIES; j++) begin
253             for (int grp = 0; grp < N_GROUPS; grp++) begin
254                 issue_one_compute(grp[1:0],
255                                 make_lane_mask(tile_base, grp, j),
256                                 px[j], py[j], m[j]);
257                 cycles++;
258                 if (cycles > RUN_TIMEOUT_CYCLES) begin
259                     $fatal(1, "ERROR: timeout while streaming wrapper computes");
260                 end
261             end
262         end
263
264         drain_tile();
265         for (int grp = 0; grp < N_GROUPS; grp++) begin
266             @(negedge clk);
267             drive_idle();
268             i_grp_sel = grp[1:0];
269             #1;
270             save_group(tile_base, grp);
271         end
272     end
273 endtask
274
275 task automatic write_output(input string fname);
276     begin
277         fo = $fopen(fname, "w");
278         if (fo == 0) $fatal(1, "ERROR: cannot open OUT_FILE=%s", fname);
279
280         $fwrite(fo, "# i ax ay (S1E8M18 hex)\n");
281         for (int i = 0; i < N_BODIES; i++) begin
282             $fwrite(fo, "%4d %07h %07h\n", i, ax[i], ay[i]);
283         end
284
285         $fclose(fo);
286     end
287 endtask
288
289 initial begin
290     rst_n = 1'b0;
291     i_grp_sel = 2'd0;
292     drive_idle();
293     err_count = 0;
294     cycles = 0;
295
296     read_frame_file(INPUT_FILE);
297
298     repeat (5) @(posedge clk);
299     @(negedge clk);
300     rst_n = 1'b1;
301     repeat (5) @(posedge clk);
302
303     for (int tile_base = 0; tile_base < N_BODIES; tile_base += TILE_SIZE) begin
304         run_tile(tile_base);
305     end
306
307     write_output(OUT_FILE);

```

```

308
309     if (err_count == 0) begin
310         $display("PASS: tb_four_core_wrapper completed without errors");
311     end else begin
312         $display("FAIL: tb_four_core_wrapper saw %0d error(s)", err_count);
313     end
314     $display("DONE. cycles=%0d wrote %s", cycles, OUT_FILE);
315     $finish;
316 end
317
318 endmodule

```

### A.4.3 tb\_core\_accel.sv

Listing 35: code/Hardware/tb/tb\_core\_accel.sv

```

1  `timescale 1ns/1ps
2
3  //-----
4  // TB: multi-body streaming scheduler driving a 2-body core
5  // - Reads frame0 file:  idx px py vx vy m   (all hex, S1E8M18 for px/py/vx/vy/m)
6  // - Streams all pairs (b1 accumulates contributions from every b2!=b1)
7  // - Core interface:
8  //     inputs: 27-bit S1E8M18 (pos/mass), prev accel: 27-bit S1E8M18
9  //     outputs: 27-bit S1E8M18
10 // - Core has fixed epsilon_square internally
11
12 module tb_core_accel;
13
14     // -----
15     // User knobs
16     // -----
17     localparam int N_BODIES = 1024;    // <-- set any N you want
18     localparam int PIPE_LAT = 18;
19     parameter int PREV_LAT = 16;
20
21     localparam string INPUT_FILE = "tb/frame_input/frame0_1024binit200_27bits.txt";
22     localparam string OUT_FILE_27bits = "tb/output/eps025_accel_27bits_1024binit200.txt";
23
24     logic          i_clk;
25     logic          i_rst;    // active-low reset
26
27     logic [26:0] i_b1_x, i_b1_y;
28     logic [26:0] i_b2_x, i_b2_y;
29     logic [26:0] i_m_b2;
30
31     wire [26:0] i_a_b1_x, i_a_b1_y;    // 27-bit prev accel, late-aligned
32     wire [26:0] o_a_b1_x, o_a_b1_y;    // 27-bit DUT accel out
33
34     two_body_core dut (
35         .i_clk      (i_clk),
36         .i_rst      (i_rst),
37
38         .i_b1_x     (i_b1_x),
39         .i_b1_y     (i_b1_y),
40         .i_b2_x     (i_b2_x),
41         .i_b2_y     (i_b2_y),
42         .i_m_b2     (i_m_b2),
43
44         .i_a_b1_x   (i_a_b1_x),
45         .i_a_b1_y   (i_a_b1_y),
46
47         .o_a_b1_x   (o_a_b1_x),
48         .o_a_b1_y   (o_a_b1_y)
49     );
50
51     // Clock
52     initial i_clk = 1'b0;
53     always #5 i_clk = ~i_clk; // 100MHz

```

```

54
55 // Frame storage
56 logic [26:0] px [0:N_BODIES-1];
57 logic [26:0] py [0:N_BODIES-1];
58 logic [26:0] m [0:N_BODIES-1];
59
60
61
62 // Streaming control
63 logic in_valid;
64 logic [PIPE_LAT-1:0] vsh;
65 logic out_valid;
66 assign out_valid = vsh[PIPE_LAT-1];
67
68 logic [26:0] launch_prev_x_27, launch_prev_y_27;
69 logic [26:0] prev_x_pipe [0:PREV_LAT-1];
70 logic [26:0] prev_y_pipe [0:PREV_LAT-1];
71
72 always_ff @(posedge i_clk) begin
73     if (!i_rst) begin
74         vsh <= '0;
75         for (int p = 0; p < PREV_LAT; p++) begin
76             prev_x_pipe[p] <= 27'd0;
77             prev_y_pipe[p] <= 27'd0;
78         end
79     end else begin
80         vsh <= {vsh[PIPE_LAT-2:0], in_valid};
81
82         prev_x_pipe[0] <= in_valid ? launch_prev_x_27 : 27'd0;
83         prev_y_pipe[0] <= in_valid ? launch_prev_y_27 : 27'd0;
84         for (int p = 0; p < PREV_LAT-1; p++) begin
85             prev_x_pipe[p+1] <= prev_x_pipe[p];
86             prev_y_pipe[p+1] <= prev_y_pipe[p];
87         end
88     end
89 end
90
91 assign i_a_b1_x = prev_x_pipe[PREV_LAT-1];
92 assign i_a_b1_y = prev_y_pipe[PREV_LAT-1];
93
94 // Helper tasks for driving input and reading frame files
95 task automatic drive_idle();
96     begin
97         i_b1_x <= 27'h1FC0000;
98         i_b1_y <= 27'h00000000;
99         i_b2_x <= 27'h20000000;
100        i_b2_y <= 27'h00000000;
101        i_m_b2 <= 27'h1FC0000;
102
103        launch_prev_x_27 <= 27'd0;
104        launch_prev_y_27 <= 27'd0;
105
106        in_valid <= 1'b0;
107    end
108 endtask
109
110 // Helper task to drive one transaction (one pair of bodies, plus prev accel)
111 task automatic drive_txn(
112     input int b1,
113     input int b2,
114     input logic [26:0] a_prev_x_27,
115     input logic [26:0] a_prev_y_27
116 );
117     begin
118         i_b1_x <= px[b1];
119         i_b1_y <= py[b1];
120         i_b2_x <= px[b2];
121         i_b2_y <= py[b2];
122         i_m_b2 <= m[b2];
123
124         launch_prev_x_27 <= a_prev_x_27;
125         launch_prev_y_27 <= a_prev_y_27;
126

```

```

127     in_valid <= 1'b1;
128     end
129 endtask
130
131 // Read input frame file
132 // line format: idx px py vx vy m (hex)
133 task automatic read_frame_file(input string fname);
134     int fd;
135     string line;
136     int idx;
137     logic [26:0] lpx, lpy, lvx, lvy, lm;
138     int got;
139     begin
140         for (int t = 0; t < N_BODIES; t++) begin
141             px[t] = 27'h0000000;
142             py[t] = 27'h0000000;
143             m[t] = 27'h0000000;
144         end
145
146         fd = $fopen(fname, "r");
147         if (fd == 0) $fatal(1, "ERROR: cannot open INPUT_FILE=%s", fname);
148
149         while (!$feof(fd)) begin
150             line = "";
151             void'($fgets(line, fd));
152
153             if (line.len() == 0) continue;
154             if (line.substr(0,0) == "#") continue;
155
156             got = $sscanf(line, "%d %h %h %h %h %h",
157                 idx, lpx, lpy, lvx, lvy, lm);
158
159             if (got == 6 && idx >= 0 && idx < N_BODIES) begin
160                 px[idx] = lpx;
161                 py[idx] = lpy;
162                 m[idx] = lm;
163             end
164         end
165         $fclose(fd);
166     end
167 endtask
168
169 // Auto-alignment queue (tracks which txn produced each output)
170 typedef struct packed {
171     int b1;
172     int b2;
173 } txn_t;
174
175 txn_t q[$];
176
177 // Scheduling state
178 // We keep per-b1 running accel in 27-bit (as the core output is "prev+term")
179 // busy[b1] prevents issuing next (b1, b2) until previous output for b1 returns.
180 logic [26:0] a_acc_x_27 [0:N_BODIES-1];
181 logic [26:0] a_acc_y_27 [0:N_BODIES-1];
182 logic        busy        [0:N_BODIES-1];
183 int         next_j       [0:N_BODIES-1];
184
185 function automatic int pick_b1();
186     int b;
187     begin
188         pick_b1 = -1;
189         for (b = 0; b < N_BODIES; b++) begin
190             if (!busy[b]) begin
191                 // skip self-pair
192                 while (next_j[b] < N_BODIES && next_j[b] == b) next_j[b]++;
193                 if (next_j[b] < N_BODIES) begin
194                     pick_b1 = b;
195                     return pick_b1;
196                 end
197             end
198         end
199     end

```

```

200 endfunction
201
202 function automatic bit all_done();
203     int b;
204     begin
205         all_done = 1'b1;
206         for (b = 0; b < N_BODIES; b++) begin
207             int tmp;
208             tmp = next_j[b];
209             while (tmp < N_BODIES && tmp == b) tmp++;
210             if (tmp < N_BODIES) all_done = 1'b0;
211             if (busy[b]) all_done = 1'b0;
212         end
213         if (q.size() != 0) all_done = 1'b0;
214         if (vsh != '0) all_done = 1'b0;
215     end
216 endfunction
217
218 integer fo;
219 integer b;
220 integer b1;
221 integer b2;
222
223 initial begin
224     drive_idle();
225
226     // reset (active-low)
227     i_rst = 1'b0;
228     repeat (5) @(posedge i_clk);
229     i_rst = 1'b1;
230     repeat (5) @(posedge i_clk);
231
232     read_frame_file(INPUT_FILE);
233
234     // init per-body accumulators and sched
235     for (b = 0; b < N_BODIES; b++) begin
236         a_acc_x_27[b] = 27'd0;
237         a_acc_y_27[b] = 27'd0;
238         busy[b] = 1'b0;
239         next_j[b] = 0;
240     end
241
242     // streaming loop
243     while (!all_done()) begin
244         @(posedge i_clk); #1;
245
246         // 1) receive aligned output (27-bit), update per-b1 accumulator
247         if (out_valid) begin
248             txn_t t;
249             if (q.size() == 0) $fatal(1, "out_valid but txn queue empty (alignment bug)");
250             t = q.pop_front();
251
252             // core returns prev+term already, so we overwrite accumulator
253             a_acc_x_27[t.b1] = o_a_b1_x;
254             a_acc_y_27[t.b1] = o_a_b1_y;
255
256             busy[t.b1] = 1'b0;
257         end
258
259         // 2) issue one txn if possible
260         b1 = pick_b1();
261         if (b1 >= 0) begin
262             b2 = next_j[b1];
263             while (b2 < N_BODIES && b2 == b1) b2++;
264
265             if (b2 < N_BODIES) begin
266                 drive_txn(b1, b2, a_acc_x_27[b1], a_acc_y_27[b1]);
267
268                 q.push_back('{b1:b1, b2:b2});
269
270                 busy[b1] = 1'b1;
271                 next_j[b1] = b2 + 1;
272             end else begin

```

```

273     drive_idle();
274     end
275 end else begin
276     drive_idle();
277 end
278 end
279
280 // write 27-bit accumulated output
281 fo = $fopen(OUT_FILE_27bits, "w");
282 if (fo == 0) $fatal(1, "ERROR: cannot open OUT_FILE_27bits=%s", OUT_FILE_27bits);
283 $fwrite(fo, "# i ax27 ay27 (S1E8M18 hex)\n");
284 for (b = 0; b < N_BODIES; b++) begin
285     $fwrite(fo, "%4d %07h %07h\n", b, a_acc_x_27[b], a_acc_y_27[b]);
286 end
287 $fclose(fo);
288
289 $display("DONE. Wrote %s", OUT_FILE_27bits);
290 $finish;
291 end
292
293 endmodule

```

#### A.4.4 tb\_nbody\_integrator.sv

Listing 36: code/Hardware/tb/tb\_nbody\_integrator.sv

```

1  `timescale 1ns/1ps
2
3  module tb_nbody_integrator;
4
5      localparam int DATA_W = 27;
6      localparam int SIGN_SHIFT = 26;
7
8      logic clk;
9      logic reset;
10     logic i_start;
11     logic o_done;
12
13     logic [DATA_W-1:0] i_x;
14     logic [DATA_W-1:0] i_y;
15     logic [DATA_W-1:0] i_vx;
16     logic [DATA_W-1:0] i_vy;
17     logic [DATA_W-1:0] i_ax;
18     logic [DATA_W-1:0] i_ay;
19
20     logic [DATA_W-1:0] o_x;
21     logic [DATA_W-1:0] o_y;
22     logic [DATA_W-1:0] o_vx;
23     logic [DATA_W-1:0] o_vy;
24
25     int err_count;
26
27     initial clk = 1'b0;
28     always #5 clk = ~clk;
29
30     nbody_integrator #(
31         .DATA_W(DATA_W)
32     ) dut (
33         .clk      (clk),
34         .reset    (reset),
35         .i_start  (i_start),
36         .o_done   (o_done),
37         .i_x      (i_x),
38         .i_y      (i_y),
39         .i_vx     (i_vx),
40         .i_vy     (i_vy),
41         .i_ax     (i_ax),
42         .i_ay     (i_ay),
43         .o_x      (o_x),

```

```

44     .o_y    (o_y),
45     .o_vx  (o_vx),
46     .o_vy  (o_vy)
47 );
48
49 function automatic logic [17:0] rtl_frac18_from_u27(input logic [DATA_W-1:0] u);
50     begin
51         rtl_frac18_from_u27 = {1'b1, u[17:1]};
52     end
53 endfunction
54
55 function automatic int lod37_shift_amt(input logic [36:0] x);
56     begin
57         lod37_shift_amt = 37;
58         for (int bit_idx = 36; bit_idx >= 0; bit_idx--) begin
59             if (x[bit_idx] && lod37_shift_amt == 37) begin
60                 lod37_shift_amt = 36 - bit_idx;
61             end
62         end
63     end
64 endfunction
65
66 function automatic logic [DATA_W-1:0] fp27_add_model(
67     input logic [DATA_W-1:0] a,
68     input logic [DATA_W-1:0] b
69 );
70     logic sa, sb;
71     int ea, eb;
72     logic [17:0] af, bf;
73     logic a_larger;
74     int exp_diff_a, exp_diff_b;
75     int larger_exp;
76     logic [36:0] a_ext, b_ext;
77     logic [36:0] a_shifted, b_shifted;
78     logic [36:0] pre_sum;
79     int shft_amt;
80     logic [53:0] pre_frac_54;
81     logic [53:0] pre_frac_shft;
82     logic [53:0] uflow_shift;
83     logic [17:0] osum_f;
84     int osum_e;
85     logic underflow;
86     logic out_sign;
87     begin
88         sa = a[SIGN_SHIFT];
89         sb = b[SIGN_SHIFT];
90         ea = a[25:18];
91         eb = b[25:18];
92         af = rtl_frac18_from_u27(a);
93         bf = rtl_frac18_from_u27(b);
94
95         a_larger = ((ea > eb) || ((ea == eb) && (af > bf)));
96         exp_diff_a = (eb - ea) & 8'hff;
97         exp_diff_b = (ea - eb) & 8'hff;
98         larger_exp = (eb > ea) ? eb : ea;
99
100        a_ext = {1'b0, af, 18'b0};
101        b_ext = {1'b0, bf, 18'b0};
102        a_shifted = a_larger ? a_ext : ((exp_diff_a > 35) ? 37'd0 : (a_ext >> exp_diff_a));
103        b_shifted = !a_larger ? b_ext : ((exp_diff_b > 35) ? 37'd0 : (b_ext >> exp_diff_b));
104
105        if ((sa ^ sb) && a_larger) begin
106            pre_sum = a_shifted - b_shifted;
107        end else if ((sa ^ sb) && !a_larger) begin
108            pre_sum = b_shifted - a_shifted;
109        end else begin
110            pre_sum = a_shifted + b_shifted;
111        end
112
113        shft_amt = lod37_shift_amt(pre_sum);
114        pre_frac_54 = {pre_sum, 17'b0};
115        pre_frac_shft = pre_frac_54 << (shft_amt + 1);
116        uflow_shift = pre_frac_54 << shft_amt;

```

```

117     osum_f = pre_frac_shft[53:36];
118     osum_e = (larger_exp - shft_amt + 1) & 8'hff;
119     underflow = ~uflow_shift[53];
120     out_sign = a_larger ? sa : sb;
121
122     if ((ea == 0) && (eb == 0)) begin
123         fp27_add_model = '0;
124     end else if (ea == 0) begin
125         fp27_add_model = b;
126     end else if (eb == 0) begin
127         fp27_add_model = a;
128     end else if (underflow) begin
129         fp27_add_model = '0;
130     end else if (pre_sum == 0) begin
131         fp27_add_model = '0;
132     end else begin
133         fp27_add_model = {out_sign, osum_e[7:0], osum_f};
134     end
135 end
136 endfunction
137
138 task automatic clear_inputs;
139     begin
140         i_start = 1'b0;
141         i_x = '0;
142         i_y = '0;
143         i_vx = '0;
144         i_vy = '0;
145         i_ax = '0;
146         i_ay = '0;
147     end
148 endtask
149
150 task automatic run_case(
151     input string name,
152     input logic [DATA_W-1:0] x,
153     input logic [DATA_W-1:0] y,
154     input logic [DATA_W-1:0] vx,
155     input logic [DATA_W-1:0] vy,
156     input logic [DATA_W-1:0] ax,
157     input logic [DATA_W-1:0] ay
158 );
159     logic [DATA_W-1:0] exp_vx;
160     logic [DATA_W-1:0] exp_vy;
161     logic [DATA_W-1:0] exp_x;
162     logic [DATA_W-1:0] exp_y;
163     int wait_cycles;
164     begin
165         exp_vx = fp27_add_model(vx, ax);
166         exp_vy = fp27_add_model(vy, ay);
167         exp_x = fp27_add_model(x, exp_vx);
168         exp_y = fp27_add_model(y, exp_vy);
169
170         @(negedge clk);
171         i_x = x;
172         i_y = y;
173         i_vx = vx;
174         i_vy = vy;
175         i_ax = ax;
176         i_ay = ay;
177         i_start = 1'b1;
178         @(posedge clk);
179         @(negedge clk);
180         i_start = 1'b0;
181
182         wait_cycles = 0;
183         do begin
184             @(posedge clk);
185             #1;
186             wait_cycles++;
187         end while (o_done != 1'b1 && wait_cycles < 20);
188
189         if (o_done != 1'b1) begin

```

```

190     $display("ERROR %s timed out waiting for o_done", name);
191     err_count++;
192 end else if (wait_cycles != 9) begin
193     $display("ERROR %s done latency got %0d expected 9", name, wait_cycles);
194     err_count++;
195 end
196
197 if (o_vx != exp_vx || o_vy != exp_vy || o_x != exp_x || o_y != exp_y) begin
198     $display("ERROR %s mismatch", name);
199     $display(" got vx=%07h vy=%07h x=%07h y=%07h", o_vx, o_vy, o_x, o_y);
200     $display(" exp vx=%07h vy=%07h x=%07h y=%07h", exp_vx, exp_vy, exp_x, exp_y);
201     err_count++;
202 end else begin
203     $display("PASS %s vx=%07h vy=%07h x=%07h y=%07h", name, o_vx, o_vy, o_x, o_y);
204 end
205
206 @(posedge clk);
207 #1;
208 if (o_done != 1'b0) begin
209     $display("ERROR %s o_done did not deassert after one cycle", name);
210     err_count++;
211 end
212 end
213 endtask
214
215 initial begin
216     err_count = 0;
217     reset = 1'b1;
218     clear_inputs();
219
220 repeat (3) @(posedge clk);
221 #1;
222 if (o_done != 1'b0 || o_x != '0 || o_y != '0 || o_vx != '0 || o_vy != '0) begin
223     $display("ERROR reset outputs not zero");
224     err_count++;
225 end
226
227 @(negedge clk);
228 reset = 1'b0;
229
230 run_case("zero", 27'h0000000, 27'h0000000, 27'h0000000, 27'h0000000,
231         27'h0000000, 27'h0000000);
232
233 run_case("positive", 27'h3fc0000, 27'h4040000, 27'h3f80000, 27'h3f00000,
234         27'h3e80000, 27'h3e00000);
235
236 run_case("signed", 27'h4040000, 27'h6040000, 27'h3fc0000, 27'h5fc0000,
237         27'h5fc0000, 27'h3fc0000);
238
239 run_case("mantissa", 27'h3fc1234, 27'h4045678, 27'h3f8abcd, 27'h5f81234,
240         27'h3e85555, 27'h3e8aaaa);
241
242 if (err_count == 0) begin
243     $display("PASS: tb_nbody_integrator completed without errors");
244 end else begin
245     $display("FAIL: tb_nbody_integrator saw %0d error(s)", err_count);
246 end
247
248 $finish;
249 end
250
251 endmodule

```

#### A.4.5 tb\_nbody\_mem.sv

Listing 37: code/Hardware/tb/tb\_nbody\_mem.sv

```

1 `timescale 1ns/1ps
2

```

```

3 module tb_nbody_mem;
4
5     localparam int MAX_BODIES = 16;
6     localparam int DATA_W = 27;
7     localparam int PTR_W = 4;
8
9     logic clk;
10
11     logic          cpu_body_we;
12     logic [PTR_W-1:0] cpu_body_waddr;
13     logic [DATA_W-1:0] cpu_x;
14     logic [DATA_W-1:0] cpu_y;
15     logic [DATA_W-1:0] cpu_m;
16     logic [DATA_W-1:0] cpu_vx;
17     logic [DATA_W-1:0] cpu_vy;
18
19     logic [PTR_W-1:0] body_raddr;
20     logic [DATA_W-1:0] body_x;
21     logic [DATA_W-1:0] body_y;
22     logic [DATA_W-1:0] body_m;
23     logic [DATA_W-1:0] body_vx;
24     logic [DATA_W-1:0] body_vy;
25     logic [DATA_W-1:0] body_ax;
26     logic [DATA_W-1:0] body_ay;
27
28     logic          body_update_we;
29     logic [PTR_W-1:0] body_update_addr;
30     logic [DATA_W-1:0] body_update_x;
31     logic [DATA_W-1:0] body_update_y;
32     logic [DATA_W-1:0] body_update_vx;
33     logic [DATA_W-1:0] body_update_vy;
34
35     logic          accel_we;
36     logic [PTR_W-1:0] accel_waddr;
37     logic [DATA_W-1:0] accel_ax;
38     logic [DATA_W-1:0] accel_ay;
39
40     int err_count;
41
42     initial clk = 1'b0;
43     always #5 clk = ~clk;
44
45     nbody_mem #(
46         .MAX_BODIES(MAX_BODIES),
47         .DATA_W(DATA_W),
48         .PTR_W(PTR_W)
49     ) dut (
50         .clk          (clk),
51         .cpu_body_we  (cpu_body_we),
52         .cpu_body_waddr (cpu_body_waddr),
53         .cpu_x        (cpu_x),
54         .cpu_y        (cpu_y),
55         .cpu_m        (cpu_m),
56         .cpu_vx       (cpu_vx),
57         .cpu_vy       (cpu_vy),
58         .body_raddr   (body_raddr),
59         .body_x       (body_x),
60         .body_y       (body_y),
61         .body_m       (body_m),
62         .body_vx      (body_vx),
63         .body_vy      (body_vy),
64         .body_ax      (body_ax),
65         .body_ay      (body_ay),
66         .body_update_we (body_update_we),
67         .body_update_addr (body_update_addr),
68         .body_update_x (body_update_x),
69         .body_update_y (body_update_y),
70         .body_update_vx (body_update_vx),
71         .body_update_vy (body_update_vy),
72         .accel_we      (accel_we),
73         .accel_waddr   (accel_waddr),
74         .accel_ax      (accel_ax),
75         .accel_ay      (accel_ay)

```

```

76 );
77
78 task automatic clear_inputs;
79 begin
80     cpu_body_we      = 1'b0;
81     cpu_body_waddr   = '0;
82     cpu_x            = '0;
83     cpu_y            = '0;
84     cpu_m            = '0;
85     cpu_vx           = '0;
86     cpu_vy           = '0;
87     body_raddr       = '0;
88     body_update_we   = 1'b0;
89     body_update_addr = '0;
90     body_update_x    = '0;
91     body_update_y    = '0;
92     body_update_vx   = '0;
93     body_update_vy   = '0;
94     accel_we         = 1'b0;
95     accel_waddr      = '0;
96     accel_ax         = '0;
97     accel_ay         = '0;
98 end
99 endtask
100
101 task automatic cpu_write_body(
102     input [PTR_W-1:0] addr,
103     input [DATA_W-1:0] x,
104     input [DATA_W-1:0] y,
105     input [DATA_W-1:0] mass,
106     input [DATA_W-1:0] vx,
107     input [DATA_W-1:0] vy
108 );
109 begin
110     @(negedge clk);
111     cpu_body_we      = 1'b1;
112     cpu_body_waddr   = addr;
113     cpu_x            = x;
114     cpu_y            = y;
115     cpu_m            = mass;
116     cpu_vx           = vx;
117     cpu_vy           = vy;
118     @(posedge clk);
119     @(negedge clk);
120     cpu_body_we      = 1'b0;
121 end
122 endtask
123
124 task automatic update_body(
125     input [PTR_W-1:0] addr,
126     input [DATA_W-1:0] x,
127     input [DATA_W-1:0] y,
128     input [DATA_W-1:0] vx,
129     input [DATA_W-1:0] vy
130 );
131 begin
132     @(negedge clk);
133     body_update_we   = 1'b1;
134     body_update_addr = addr;
135     body_update_x    = x;
136     body_update_y    = y;
137     body_update_vx   = vx;
138     body_update_vy   = vy;
139     @(posedge clk);
140     @(negedge clk);
141     body_update_we   = 1'b0;
142 end
143 endtask
144
145 task automatic write_accel(
146     input [PTR_W-1:0] addr,
147     input [DATA_W-1:0] ax,
148     input [DATA_W-1:0] ay

```

```

149 );
150   begin
151     @(negedge clk);
152     accel_we      = 1'b1;
153     accel_waddr   = addr;
154     accel_ax      = ax;
155     accel_ay      = ay;
156     @(posedge clk);
157     @(negedge clk);
158     accel_we      = 1'b0;
159   end
160 endtask
161
162 task automatic check_body(
163   input [PTR_W-1:0] addr,
164   input [DATA_W-1:0] exp_x,
165   input [DATA_W-1:0] exp_y,
166   input [DATA_W-1:0] exp_m,
167   input [DATA_W-1:0] exp_vx,
168   input [DATA_W-1:0] exp_vy,
169   input [DATA_W-1:0] exp_ax,
170   input [DATA_W-1:0] exp_ay
171 );
172   begin
173     @(negedge clk);
174     body_raddr = addr;
175     @(posedge clk);
176     #1;
177     if (body_x != exp_x || body_y != exp_y || body_m != exp_m ||
178         body_vx != exp_vx || body_vy != exp_vy ||
179         body_ax != exp_ax || body_ay != exp_ay) begin
180       $display("ERROR addr=%0d got x=%07h y=%07h m=%07h vx=%07h vy=%07h ax=%07h ay=%07h",
181         addr, body_x, body_y, body_m, body_vx, body_vy, body_ax, body_ay);
182       $display("          exp x=%07h y=%07h m=%07h vx=%07h vy=%07h ax=%07h ay=%07h",
183         exp_x, exp_y, exp_m, exp_vx, exp_vy, exp_ax, exp_ay);
184       err_count++;
185     end
186   end
187 endtask
188
189 task automatic simultaneous_cpu_update_accel;
190   begin
191     @(negedge clk);
192     cpu_body_we      = 1'b1;
193     cpu_body_waddr   = 4'd7;
194     cpu_x            = 27'h11111111;
195     cpu_y            = 27'h12222222;
196     cpu_m            = 27'h13333333;
197     cpu_vx           = 27'h14444444;
198     cpu_vy           = 27'h15555555;
199
200     body_update_we   = 1'b1;
201     body_update_addr = 4'd7;
202     body_update_x    = 27'h21111111;
203     body_update_y    = 27'h22222222;
204     body_update_vx   = 27'h24444444;
205     body_update_vy   = 27'h25555555;
206
207     accel_we         = 1'b1;
208     accel_waddr      = 4'd7;
209     accel_ax         = 27'h26666666;
210     accel_ay         = 27'h27777777;
211
212     @(posedge clk);
213     @(negedge clk);
214     cpu_body_we      = 1'b0;
215     body_update_we   = 1'b0;
216     accel_we         = 1'b0;
217   end
218 endtask
219
220 initial begin
221   err_count = 0;

```

```

222 clear_inputs();
223 repeat (2) @(posedge clk);
224
225 cpu_write_body(4'd3, 27'h1010001, 27'h1020002, 27'h1030003, 27'h1040004, 27'h1050005);
226 check_body(4'd3, 27'h1010001, 27'h1020002, 27'h1030003, 27'h1040004, 27'h1050005,
227           27'h0000000, 27'h0000000);
228
229 write_accel(4'd3, 27'h1a60006, 27'h1a70007);
230 check_body(4'd3, 27'h1010001, 27'h1020002, 27'h1030003, 27'h1040004, 27'h1050005,
231           27'h1a60006, 27'h1a70007);
232
233 update_body(4'd3, 27'h2010001, 27'h2020002, 27'h2040004, 27'h2050005);
234 check_body(4'd3, 27'h2010001, 27'h2020002, 27'h1030003, 27'h2040004, 27'h2050005,
235           27'h1a60006, 27'h1a70007);
236
237 cpu_write_body(4'd5, 27'h3010001, 27'h3020002, 27'h3030003, 27'h3040004, 27'h3050005);
238 check_body(4'd5, 27'h3010001, 27'h3020002, 27'h3030003, 27'h3040004, 27'h3050005,
239           27'h0000000, 27'h0000000);
240 check_body(4'd3, 27'h2010001, 27'h2020002, 27'h1030003, 27'h2040004, 27'h2050005,
241           27'h1a60006, 27'h1a70007);
242
243 simultaneous_cpu_update_accel();
244 check_body(4'd7, 27'h2111111, 27'h2222222, 27'h1333333, 27'h2444444, 27'h2555555,
245           27'h2666666, 27'h2777777);
246
247 if (err_count == 0) begin
248     $display("PASS: tb_nbody_mem completed without errors");
249 end else begin
250     $display("FAIL: tb_nbody_mem saw %0d error(s)", err_count);
251 end
252
253 $finish;
254 end
255
256 endmodule

```

#### A.4.6 tb\_fpadd.sv

Listing 38: code/Hardware/tb/tb\_fpadd.sv

```

1  `timescale 1ns/1ps
2
3  module tb_fpadd;
4
5      reg        iCLK;
6      reg [26:0] iA;
7      reg [26:0] iB;
8      wire [26:0] oSum;
9      reg [26:0] exp_out;
10
11     integer fd_in;
12     integer ret;
13     integer case_idx;
14     integer pass_cnt;
15     integer fail_cnt;
16
17     FpAdd dut (
18         .iCLK(iCLK),
19         .iA (iA),
20         .iB (iB),
21         .oSum(oSum)
22     );
23
24     initial begin
25         iCLK = 0;
26         forever #5 iCLK = ~iCLK;
27     end
28
29     initial begin

```

```

30     iA = 0;
31     iB = 0;
32     exp_out = 0;
33     case_idx = 0;
34     pass_cnt = 0;
35     fail_cnt = 0;
36
37     fd_in = $fopen("tb/frame_input/fpadd_cases.txt", "r");
38     if (fd_in == 0) begin
39         $display("ERROR: cannot open fpadd_cases.txt");
40         $finish;
41     end
42
43     @(negedge iCLK);
44
45     while (!$feof(fd_in)) begin
46         ret = $fscanf(fd_in, "%h %h %h\n", iA, iB, exp_out);
47
48         if (ret == 3) begin
49             // FpAdd is two-stage pipelined, wait 2 cycles for output to settle
50             @(posedge iCLK);
51             @(posedge iCLK);
52             #1;
53
54             if (oSum === exp_out) begin
55                 pass_cnt = pass_cnt + 1;
56                 $display("[%0d] a=%07h b=%07h rtl=%07h py=%07h PASS",
57                     case_idx, iA, iB, oSum, exp_out);
58             end
59             else begin
60                 fail_cnt = fail_cnt + 1;
61                 $display("[%0d] a=%07h b=%07h rtl=%07h py=%07h FAIL",
62                     case_idx, iA, iB, oSum, exp_out);
63                 // Uncomment below to stop on the first failure
64                 // $stop;
65             end
66
67             case_idx = case_idx + 1;
68
69             // Load next input group with a clean negedge clock edge
70             @(negedge iCLK);
71         end
72         else begin
73             ret = $fgetc(fd_in);
74         end
75     end
76
77     $display("=====");
78     $display("tb_fpadd done");
79     $display("total = %0d, pass = %0d, fail = %0d",
80         case_idx, pass_cnt, fail_cnt);
81     $display("=====");
82
83     $fclose(fd_in);
84     $finish;
85 end
86
87 endmodule

```

#### A.4.7 tb\_fpmul.sv

Listing 39: code/Hardware/tb/tb\_fpmul.sv

```

1  `timescale 1ns/1ps
2
3  module tb_fpmul;
4
5      reg [26:0] iA;
6      reg [26:0] iB;

```

```

7   wire [26:0] oProd;
8   reg  [26:0] exp_out;
9
10  integer fd_in;
11  integer ret;
12  integer case_idx;
13
14  FpMul dut (
15      .iA(iA),
16      .iB(iB),
17      .oProd(oProd)
18  );
19
20  initial begin
21      iA = 0;
22      iB = 0;
23      exp_out = 0;
24
25      fd_in = $fopen("tb/frame_input/fpmul_cases.txt", "r");
26      if (fd_in == 0) begin
27          $display("ERROR: cannot open fpmul_cases.txt");
28          $finish;
29      end
30
31      case_idx = 0;
32
33      while (!$feof(fd_in)) begin
34          ret = $fscanf(fd_in, "%h %h %h\n", iA, iB, exp_out);
35          $display("ret=%0d iA=%h iB=%h exp=%h", ret, iA, iB, exp_out);
36
37          if (ret == 3) begin
38              #1;
39              $display("[%0d] rtl=%h py=%h %s",
40                  case_idx, oProd, exp_out,
41                  (oProd == exp_out) ? "PASS" : "FAIL");
42              case_idx = case_idx + 1;
43          end
44          else begin
45              // consume one bad line if needed
46              ret = $fgetc(fd_in);
47          end
48      end
49
50      $fclose(fd_in);
51      $finish;
52  end
53
54  endmodule

```

#### A.4.8 tb\_fpinvsqrt.sv

Listing 40: code/Hardware/tb/tb\_fpinvsqrt.sv

```

1  `timescale 1ns/1ps
2
3  module tb_fpinvsqrt;
4
5      reg          iCLK;
6      reg  [26:0] iA;
7      wire [26:0] oInvSqrt;
8      reg  [26:0] exp_out;
9
10     integer fd_in;
11     integer ret;
12     integer case_idx;
13     integer pass_cnt;
14     integer fail_cnt;
15     integer k;
16

```

```

17 localparam integer SETTLE_CYCLES = 8;
18
19 FpInvSqrt dut (
20     .iCLK      (iCLK),
21     .iA        (iA),
22     .oInvSqrt (oInvSqrt)
23 );
24
25 initial begin
26     iCLK = 1'b0;
27     forever #5 iCLK = ~iCLK;
28 end
29
30 initial begin
31     iA = 27'd0;
32     exp_out = 27'd0;
33     case_idx = 0;
34     pass_cnt = 0;
35     fail_cnt = 0;
36
37     fd_in = $fopen("tb/frame_input/fpinvsqrt_cases.txt", "r");
38     if (fd_in == 0) begin
39         $display("ERROR: cannot open fpinvsqrt_cases.txt");
40         $finish;
41     end
42
43     repeat (3) @(posedge iCLK); // let pipeline settle
44
45     while (!$feof(fd_in)) begin
46         ret = $fscanf(fd_in, "%h %h\n", iA, exp_out);
47
48         if (ret == 2) begin
49             // hold this input constant long enough
50             for (k = 0; k < SETTLE_CYCLES; k = k + 1)
51                 @(posedge iCLK);
52             #1;
53
54             if (oInvSqrt === exp_out) begin
55                 pass_cnt = pass_cnt + 1;
56                 $display("[%0d] x=%07h rtl=%07h py=%07h PASS",
57                     case_idx, iA, oInvSqrt, exp_out);
58             end
59             else begin
60                 fail_cnt = fail_cnt + 1;
61                 $display("[%0d] x=%07h rtl=%07h py=%07h FAIL",
62                     case_idx, iA, oInvSqrt, exp_out);
63
64                 // first-fail stop is very useful here
65                 $display("---- first fail detail ----");
66                 $display("x      = %07h", iA);
67                 $display("rtl   = %07h", oInvSqrt);
68                 $display("python = %07h", exp_out);
69                 $stop;
70             end
71
72             case_idx = case_idx + 1;
73         end
74         else begin
75             ret = $fgetc(fd_in);
76         end
77     end
78
79     $display("=====");
80     $display("tb_fpinvsqrt done");
81     $display("total = %0d, pass = %0d, fail = %0d",
82         case_idx, pass_cnt, fail_cnt);
83     $display("=====");
84
85     $fclose(fd_in);
86     $finish;
87 end
88
89 endmodule

```

## A.5 Golden Model and Input Generation Code

### A.5.1 golden.py

Listing 41: code/Golden/golden.py

```
1  # ## 3. Python Model for FP modules
2
3  # ### 3.1 FP27 Add
4  # fp27_add_model.py
5  # Standalone RTL-mimic model for FpAdd (S1E8M18 core, 27-bit)
6
7  import random
8
9  # =====
10 # format constants
11 # =====
12 U27_MASK = (1 << 27) - 1
13 SIGN_SHIFT = 26
14 EXP_SHIFT = 18
15 E_MASK = 0xFF
16 M18_MASK = (1 << 18) - 1
17
18
19 def u27_pack(sign: int, exp: int, mant: int) -> int:
20     return ((sign & 1) << SIGN_SHIFT) | ((exp & E_MASK) << EXP_SHIFT) | (mant & M18_MASK)
21
22
23 def u27_fields(u: int):
24     u &= U27_MASK
25     s = (u >> SIGN_SHIFT) & 1
26     e = (u >> EXP_SHIFT) & E_MASK
27     m = u & M18_MASK
28     return s, e, m
29
30
31 def u27_hex(u: int) -> str:
32     return f"{u & U27_MASK:07X}"
33
34
35 def rtl_frac18_from_u27(u: int) -> int:
36     """
37     Match RTL:
38     assign A_f = {1'b1, iA[17:1]};
39     => internal 18-bit mantissa used by RTL add
40     """
41     return (1 << 17) | ((u >> 1) & 0x1FFFF)
42
43
44 def lod37_shift_amt(x: int) -> int:
45     """
46     Match RTL shift_amt:
47     pre_frac[36] ? 0 :
48     pre_frac[35] ? 1 :
49     ...
50     pre_frac[0] ? 36 :
51                 37;
52     """
53     x &= (1 << 37) - 1
54     if x == 0:
55         return 37
56     for bit in range(36, -1, -1):
57         if (x >> bit) & 1:
58             return 36 - bit
59     return 37
60
61
62 def fp27_add_rtl(a: int, b: int, debug: bool = False):
63     """
64     Bit-exact mimic of RTL FpAdd.
65     Functionally collapses 2 pipeline stages into one Python call.
```

```

66     ""
67     a &= U27_MASK
68     b &= U27_MASK
69
70     sa, ea, _ = u27_fields(a)
71     sb, eb, _ = u27_fields(b)
72
73     A_f = rtl_frac18_from_u27(a) # 18-bit
74     B_f = rtl_frac18_from_u27(b) # 18-bit
75
76     # -----
77     # stage 1 combinational
78     # -----
79     A_larger = 1 if ((ea > eb) or ((ea == eb) and (A_f > B_f))) else 0
80
81     exp_diff_A = (int(eb) - int(ea)) & 0xFF # 9-bit-ish spirit
82     exp_diff_B = (int(ea) - int(eb)) & 0xFF
83     larger_exp = eb if (eb > ea) else ea
84
85     # RTL:
86     # {1'b0, A_f, 18'b0} -> 37 bits
87     A_ext = (A_f << 18) & ((1 << 37) - 1)
88     B_ext = (B_f << 18) & ((1 << 37) - 1)
89
90     if A_larger:
91         A_f_shifted = A_ext
92     else:
93         A_f_shifted = 0 if exp_diff_A > 35 else (A_ext >> exp_diff_A)
94
95     if not A_larger:
96         B_f_shifted = B_ext
97     else:
98         B_f_shifted = 0 if exp_diff_B > 35 else (B_ext >> exp_diff_B)
99
100    if ((sa ^ sb) & A_larger):
101        pre_sum = (A_f_shifted - B_f_shifted) & ((1 << 37) - 1)
102    elif ((sa ^ sb) & (1 - A_larger)):
103        pre_sum = (B_f_shifted - A_f_shifted) & ((1 << 37) - 1)
104    else:
105        pre_sum = (A_f_shifted + B_f_shifted) & ((1 << 37) - 1)
106
107    # -----
108    # stage 1 registers
109    # -----
110    buf_pre_sum      = pre_sum
111    buf_larger_exp   = larger_exp
112    buf_A_e_zero     = (ea == 0)
113    buf_B_e_zero     = (eb == 0)
114    buf_A            = a
115    buf_B            = b
116    buf_oSum_s      = sa if A_larger else sb
117
118    # -----
119    # stage 2 combinational / output logic
120    # -----
121    pre_frac = buf_pre_sum
122    shft_amt = lod37_shift_amt(pre_frac)
123
124    # RTL:
125    # pre_frac_shft = {pre_frac,17'b0} << (shft_amt+1)
126    # uflow_shift   = {pre_frac,17'b0} << (shft_amt)
127    pre_frac_54    = (pre_frac << 17) & ((1 << 54) - 1)
128    pre_frac_shft  = (pre_frac_54 << (shft_amt + 1)) & ((1 << 54) - 1)
129    uflow_shift    = (pre_frac_54 << shft_amt) & ((1 << 54) - 1)
130
131    oSum_f = (pre_frac_shft >> 36) & M18_MASK
132    oSum_e = (int(buf_larger_exp) - int(shft_amt) + 1) & 0xFF
133
134    # RTL:
135    # assign underflow = ~uflow_shift[53];
136    underflow = 1 if (((uflow_shift >> 53) & 1) == 0) else 0
137
138    if buf_A_e_zero and buf_B_e_zero:

```

```

139     out = 0
140 elif buf_A_e_zero:
141     out = buf_B
142 elif buf_B_e_zero:
143     out = buf_A
144 elif underflow:
145     out = 0
146 elif pre_frac == 0:
147     out = 0
148 else:
149     out = u27_pack(buf_oSum_s, oSum_e, oSum_f)
150
151 if debug:
152     return {
153         "a": a,
154         "b": b,
155         "sa": sa,
156         "sb": sb,
157         "ea": ea,
158         "eb": eb,
159         "A_f": A_f,
160         "B_f": B_f,
161         "A_larger": A_larger,
162         "exp_diff_A": exp_diff_A,
163         "exp_diff_B": exp_diff_B,
164         "larger_exp": larger_exp,
165         "A_f_shifted": A_f_shifted,
166         "B_f_shifted": B_f_shifted,
167         "pre_sum": pre_sum,
168         "buf_pre_sum": buf_pre_sum,
169         "pre_frac": pre_frac,
170         "shft_amt": shft_amt,
171         "pre_frac_54": pre_frac_54,
172         "pre_frac_shft": pre_frac_shft,
173         "uflow_shift": uflow_shift,
174         "oSum_f": oSum_f,
175         "oSum_e": oSum_e,
176         "underflow": underflow,
177         "buf_oSum_s": buf_oSum_s,
178         "out": out,
179         "out_hex": u27_hex(out),
180     }
181 return out
182
183
184 def make_case_file(path: str, n_random: int = 500):
185     cases = []
186
187     # -----
188     # directed cases
189     # -----
190     directed = [
191         # zeros
192         (0x0000000, 0x0000000),
193         (0x0000000, 0x3FC0000),
194         (0x3FC0000, 0x0000000),
195
196         # same-sign additions
197         (0x3FC0000, 0x3FC0000),
198         (0x4040000, 0x3FC0000),
199         (0x4040000, 0x4040000),
200
201         # sign-opposite / cancellation-ish
202         (0x4440000, 0xC040000),
203         (0xC040000, 0x4040000),
204         (0x4040000, 0xC040000),
205         (0x4040000, 0xC040001), # test bit0-ish weirdness
206         (0x3FC0000, 0xBFC0000), # exact cancel
207         (0x3FC0002, 0xBFC0000), # near cancel
208
209         # same exponent / different frac
210         (0x3FC0001, 0x3FC0000),
211         (0x3FC0002, 0x3FC0004),

```

```

212         (0x7E7FFFF, 0x3FC0000),
213
214         # large exponent diff
215         (0x7C40000, 0x0440000),
216         (0x0440000, 0x7C40000),
217
218         # random looking hand-picked
219         (0x11B31C6, 0x6AB377F),
220         (0x390A3CF, 0x0F3EFAC),
221         (0x4F2FC12, 0x0899422),
222         (0x2C8993B, 0x4E9E669),
223     ]
224     cases.extend(directed)
225
226     # -----
227     # random normals / zeros
228     # -----
229     def rand_u27():
230         s = random.getrandbits(1)
231         e = 0 if random.random() < 0.08 else random.randint(1, 254)
232         m = random.getrandbits(18)
233         return u27_pack(s, e, m)
234
235     for _ in range(n_random):
236         a = rand_u27()
237         b = rand_u27()
238         cases.append((a, b))
239
240     with open(path, "w") as f:
241         for a, b in cases:
242             y = fp27_add_rtl(a, b)
243             f.write(f"{u27_hex(a)} {u27_hex(b)} {u27_hex(y)}\n")
244
245     print(f"Wrote {len(cases)} cases to {path}")
246
247
248 # if __name__ == "__main__":
249 #     make_case_file("fpadd_cases.txt", n_random=500)
250
251
252 # ### 3.2 FP27 Mul
253 # fp27_mul_model.py
254 # Standalone RTL-mimic model for FpMul (SiE8M18 core, 27-bit)
255
256 import random
257
258 # =====
259 # format constants
260 # =====
261 U27_MASK = (1 << 27) - 1
262 SIGN_SHIFT = 26
263 EXP_SHIFT = 18
264 E_MASK = 0xFF
265 M18_MASK = (1 << 18) - 1
266
267 # NOTE:
268 # RTL uses:
269 # A_f = {1'b1, iA[17:1]}
270 # so bit0 is ignored by the multiplier datapath.
271
272
273 def u27_pack(sign: int, exp: int, mant: int) -> int:
274     return ((sign & 1) << SIGN_SHIFT) | ((exp & E_MASK) << EXP_SHIFT) | (mant & M18_MASK)
275
276
277 def u27_fields(u: int):
278     u &= U27_MASK
279     s = (u >> SIGN_SHIFT) & 1
280     e = (u >> EXP_SHIFT) & E_MASK
281     m = u & M18_MASK
282     return s, e, m
283
284

```

```

285 def u27_hex(u: int) -> str:
286     return f"{u & U27_MASK:07X}"
287
288
289 def rtl_frac18_from_u27(u: int) -> int:
290     """
291     Match RTL:
292     assign A_f = {1'b1, iA[17:1]};
293     Effective 18-bit mantissa used internally:
294     bit17 = hidden 1
295     bit16:0 = stored bits [17:1]
296     """
297     return (1 << 17) | ((u >> 1) & 0x1FFFF)
298
299
300 def fp27_mul_rtl(a: int, b: int, debug: bool = False):
301     """
302     Bit-exact mimic of RTL FpMul:
303     oProd_s = A_s ^ B_s
304     pre_prod_frac = A_f * B_f
305     pre_prod_exp = A_e + B_e
306     if pre_prod_frac[35]:
307         oProd_e = pre_prod_exp - 126
308         oProd_f = pre_prod_frac[34:17]
309     else:
310         oProd_e = pre_prod_exp - 127
311         oProd_f = pre_prod_frac[33:16]
312     underflow = pre_prod_exp < 128
313     if underflow or A_e==0 or B_e==0: out=0
314     """
315     a &= U27_MASK
316     b &= U27_MASK
317
318     sa, ea, _ = u27_fields(a)
319     sb, eb, _ = u27_fields(b)
320
321     A_f = rtl_frac18_from_u27(a)
322     B_f = rtl_frac18_from_u27(b)
323
324     oProd_s = sa ^ sb
325     pre_prod_frac = A_f * B_f # 36-bit
326     pre_prod_exp = ea + eb # 9-bit in RTL spirit
327
328     if ((pre_prod_frac >> 35) & 1) == 1:
329         oProd_e = (pre_prod_exp - 126) & 0xFF
330         oProd_f = (pre_prod_frac >> 17) & M18_MASK # [34:17]
331         path = "hi"
332     else:
333         oProd_e = (pre_prod_exp - 127) & 0xFF
334         oProd_f = (pre_prod_frac >> 16) & M18_MASK # [33:16]
335         path = "lo"
336
337     underflow = pre_prod_exp < 0x80
338
339     if underflow or (ea == 0) or (eb == 0):
340         out = 0
341     else:
342         out = u27_pack(oProd_s, oProd_e, oProd_f)
343
344     if debug:
345         return {
346             "a": a,
347             "b": b,
348             "sa": sa,
349             "sb": sb,
350             "ea": ea,
351             "eb": eb,
352             "A_f": A_f,
353             "B_f": B_f,
354             "pre_prod_frac": pre_prod_frac,
355             "pre_prod_exp": pre_prod_exp,
356             "oProd_s": oProd_s,
357             "oProd_e": oProd_e,

```

```

358         "oProd_f": oProd_f,
359         "underflow": underflow,
360         "path": path,
361         "out": out,
362         "out_hex": u27_hex(out),
363     }
364     return out
365
366
367 def make_case_file(path: str, n_random: int = 200):
368     """
369     Write:
370         a_hex b_hex expected_hex
371     """
372     cases = []
373
374     # ----- directed cases -----
375     directed = [
376         (0x0000000, 0x0000000),
377         (0x0000000, 0x3FC0000),
378         (0x3FC0000, 0x0000000),
379         (0x3FC0000, 0x3FC0000),
380         (0x4040000, 0x3FC0000),
381         (0x4040000, 0x4040000),
382         (0x43C0000, 0x3FC0000),
383         (0x7E7FFFF, 0x3FC0000),
384         (0x3FC0001, 0x3FC0001), # test bit0 ignored or not
385         (0x3FC0002, 0x3FC0002),
386         (0x7C40000, 0x0440000), # possible underflow-ish
387         (0x4440000, 0xC040000), # sign mix
388         (0xC040000, 0xC040000),
389     ]
390     cases.extend(directed)
391
392     # ----- random normals / zeros -----
393     for _ in range(n_random):
394         def rand_u27():
395             s = random.getrandbits(1)
396             # mostly normal exp, sometimes zero
397             e = 0 if random.random() < 0.08 else random.randint(1, 254)
398             m = random.getrandbits(18)
399             return u27_pack(s, e, m)
400
401         a = rand_u27()
402         b = rand_u27()
403         cases.append((a, b))
404
405     with open(path, "w") as f:
406         f.write("# a_hex b_hex expected_hex\n")
407         for a, b in cases:
408             y = fp27_mul_rtl(a, b)
409             f.write(f"{u27_hex(a)} {u27_hex(b)} {u27_hex(y)}\n")
410
411     print(f"Wrote {len(cases)} cases to {path}")
412
413
414 # if __name__ == "__main__":
415 #     make_case_file("fpmul_cases.txt", n_random=500)
416
417
418 # ### 3.3 FP27 InvSqrt
419 # fp27_invsqrt_model.py
420 import random
421
422 # -----
423 # import your already-validated standalone models
424 # -----
425 # from fp27_mul_model import fp27_mul_rtl
426 # from fp27_add_model import fp27_add_rtl
427
428 # -----
429 # format constants
430 # -----

```

```

431 U27_MASK      = (1 << 27) - 1
432 SIGN_SHIFT   = 26
433 EXP_SHIFT    = 18
434 E_MASK       = 0xFF
435 M18_MASK     = (1 << 18) - 1
436
437 # from your current model
438 MAGIC_C = 49920718 # 27'd49920718
439 ADD_K   = 33423360 # 27'd33423360
440
441
442 def u27_pack(sign: int, exp: int, mant: int) -> int:
443     return ((sign & 1) << SIGN_SHIFT) | ((exp & E_MASK) << EXP_SHIFT) | (mant & M18_MASK)
444
445
446 def u27_fields(u: int):
447     u &= U27_MASK
448     s = (u >> SIGN_SHIFT) & 1
449     e = (u >> EXP_SHIFT) & E_MASK
450     m = u & M18_MASK
451     return s, e, m
452
453
454 def u27_hex(u: int) -> str:
455     return f"{u & U27_MASK:07X}"
456
457
458 def u27_neg(u: int) -> int:
459     """Match FpNegate behavior for normal finite-style custom float."""
460     u &= U27_MASK
461     #if u == 0:
462     #    return 0
463     return u ^ (1 << SIGN_SHIFT)
464
465
466 def fast_inv_sqrt_u27_rtl(i_x: int, debug: bool = False) -> int:
467     i_x &= U27_MASK
468     w_sign, w_exp, w_man = u27_fields(i_x)
469
470     # Stage 1
471     w_s12_c = (MAGIC_C - (i_x >> 1)) & U27_MASK
472
473     # IMPORTANT:
474     # RTL does NOT flush exp==0 here.
475     # It literally does {A_s, A_e-8'd1, A_f}, so exponent wraps.
476     w_s12_b = u27_pack(w_sign, (w_exp - 1) & E_MASK, w_man)
477
478     # Stage 1 mul
479     w_s12_a = fp27_mul_rtl(w_s12_c, w_s12_c)
480
481     # Stage 2 mul
482     w_s23_a = fp27_mul_rtl(w_s12_b, w_s12_a)
483     w_s23_b = w_s12_c
484
485     # Stage 3 add: iA({~add_3_in[26], add_3_in[25:0]}), iB(ADD_K)
486     w_s45_a = fp27_add_rtl(u27_neg(w_s23_a), ADD_K)
487     w_s45_b = w_s23_b
488
489     # Final mul
490     out = fp27_mul_rtl(w_s45_b, w_s45_a)
491
492     if debug:
493         return {
494             "i_x": i_x,
495             "w_sign": w_sign,
496             "w_exp": w_exp,
497             "w_man": w_man,
498             "w_s12_c": w_s12_c,
499             "w_s12_b": w_s12_b,
500             "w_s12_a": w_s12_a,
501             "w_s23_a": w_s23_a,
502             "w_s23_b": w_s23_b,
503             "w_s45_a": w_s45_a,

```

```

504         "w_s45_b": w_s45_b,
505         "out": out,
506         "out_hex": f"{out:07X}",
507     }
508     return out
509
510
511 def make_case_file(path: str, n_random: int = 500):
512     cases = []
513
514     # -----
515     # directed cases
516     # -----
517     directed = [
518         0x0000000, # zero
519         0x3FC0000, # 1.0-ish
520         0x4040000, # 2.0-ish
521         0x3C40000, # 0.5-ish
522         0x43C0000, # 8.0-ish
523         0x44C0000,
524         0x7E7FFFF, # large normal-ish
525         0x0440000, # tiny normal-ish
526         0x3FC0001, # bit0 perturb
527         0x3FC0002,
528         0x4040001,
529         0x7C40000,
530         0x0100000,
531         0x1000000,
532         0x2000000,
533     ]
534     for x in directed:
535         cases.append(x)
536
537     # -----
538     # random normals / zeros
539     # -----
540     def rand_u27():
541         s = 0 # invsqrt input should usually be non-negative r2/r2e
542         e = 0 if random.random() < 0.08 else random.randint(1, 254)
543         m = random.getrandbits(18)
544         return u27_pack(s, e, m)
545
546     for _ in range(n_random):
547         cases.append(rand_u27())
548
549     with open(path, "w") as f:
550         for x in cases:
551             y = fast_inv_sqrt_u27_rtl(x)
552             f.write(f"{u27_hex(x)} {u27_hex(y)}\n")
553
554     print(f"Wrote {len(cases)} cases to {path}")
555
556
557 # if __name__ == "__main__":
558 #     make_case_file("fpinvsqrt_cases.txt", n_random=500)
559
560
561
562 # ## 4. Full Hardware Model
563
564 # v4real_chiplike_f187_output.py
565 # IO: 27-bit custom float S1E8M18 (sign 1, exp 8, mant 18)
566 # CORE: 27-bit custom float S1E8M18 (sign 1, exp 8, mant 18)
567 #
568 # Outputs (N_FRAMES steps, frames numbered 0,1,2,...):
569 #   OUT_ROOT/
570 #     frame_chip_input/ : chip input frames (HEX S1E8M18). frame0.txt copied from
571 #     INPUT_FILE
572 #     accel_chip_output/ : chip output accelerations (HEX S1E8M18). accel0.txt, accel1.txt
573 #     , ...
574 #     accel_core27_output/ : internal accumulated accel (HEX S1E8M18). accel27_0.txt,
575 #     accel27_1.txt, ...
576 #     frame_host_output/ : host frames (DECIMAL). frame0.txt, frame1.txt, ...

```

```

574 #
575 # Physics convention:
576 # - Chip accel is "chiplike" and MUST NOT include GO.
577 # - Host update uses:
578 #     K_HOST = GO * DT    (only used in v-update)
579 #     v += a_chip * K_HOST
580 #     p += v * DT        (DT remains here; do NOT multiply GO into this)
581 #
582 # IMPORTANT (BIT-EXACT CORE):
583 # - Core ops (add/mul/neg/fast_inv_sqrt) NEVER use float.
584 # - Core arithmetic is imported from standalone RTL-mimic models.
585 # - Host boundary quantization still uses TRUNC; denorm flush to 0.
586 # - Overflow clamps to max finite (exp=254, mant=all-1).
587
588 import os
589 import shutil
590 import numpy as np
591
592 # =====
593 # IMPORT BIT-EXACT CORE (match standalone RTL-mimic modules)
594 # =====
595 # The functions fp27_add_rtl, fp27_mul_rtl, and fast_inv_sqrt_u27_rtl
596 # are defined in earlier cells. We will assign them directly to aliases.
597 fp27_add = fp27_add_rtl
598 fp27_mul = fp27_mul_rtl
599 fast_inv_sqrt_u27 = fast_inv_sqrt_u27_rtl
600
601 # =====
602 # Top-level knobs
603 # =====
604 INPUT_FILE = "../Hardware/tb/frame_input/frame0_1024binit200_27bits.txt"
605 OUT_ROOT = "../Hardware/tb/output/eps025_v4real_chiplike_1024binit200"
606 # INPUT_FILE = "frame0_8binit10.txt"
607 # OUT_ROOT = "eps025_v4real_chiplike_8binit10"
608
609 DIR_FRAME_HOST = os.path.join(OUT_ROOT, "frame_host_output") # decimal
610 DIR_FRAME_CHIP = os.path.join(OUT_ROOT, "frame_chip_input") # hex (S1E8M18)
611 DIR_ACCEL_CHIP = os.path.join(OUT_ROOT, "accel_chip_output") # hex (S1E8M18)
612 DIR_ACCEL_CORE27 = os.path.join(OUT_ROOT, "accel_core27_output") # hex (S1E8M18)
613
614 N_FRAMES = 4
615
616 # Host integrator settings
617 DT = 0.1
618 GO = 10.0
619 K_HOST = GO * DT
620
621 HOST_DTYPE = np.float32
622
623 # EPS^2 bit-exact constant in u27.
624 # Matches RTL: localparam logic [26:0] EPSILON_SQUARE = {1'b0, 8'd125, 18'd0};
625 # This encodes 2^(125 - 127) = 0.25, so epsilon = 0.5.
626 EPS_SQUARE_U27 = (0 << 26) | (125 << 18) | 0
627
628 # =====
629 # Shared float layout params
630 # =====
631 BIAS = 127
632 E_BITS = 8
633 E_MASK = (1 << E_BITS) - 1
634
635 # =====
636 # CORE format: S1E8M18 (27-bit)
637 # =====
638 M18_BITS = 18
639 U27_MASK = (1 << 27) - 1
640 SIGN27_SHIFT = 26
641 EXP27_SHIFT = M18_BITS
642 M18_MASK = (1 << M18_BITS) - 1
643
644 def u27_fields(u27: int):
645     u27 &= U27_MASK
646     s = (u27 >> SIGN27_SHIFT) & 1

```

```

647     e = (u27 >> EXP27_SHIFT) & E_MASK
648     m = u27 & M18_MASK
649     return s, e, m
650
651 def u27_pack(s: int, e: int, m: int) -> int:
652     return ((s & 1) << SIGN27_SHIFT) | ((e & E_MASK) << EXP27_SHIFT) | (m & M18_MASK)
653
654 def u27_neg(u27: int) -> int:
655     # FIX: match RTL FpNegate exactly: {~A_s, A_e, A_f}
656     # RTL unconditionally flips bit26, no zero guard.
657     return (u27 ^ (1 << SIGN27_SHIFT)) & U27_MASK
658
659 def u27_to_hex(u27: int) -> str:
660     return f"{u27 & U27_MASK:07X}"
661
662 def u27_from_hex(h: str) -> int:
663     return int(h, 16) & U27_MASK
664
665 def u27_to_float(u27: int) -> float:
666     """For host-side view only."""
667     u27 &= U27_MASK
668     if u27 == 0:
669         return 0.0
670     s, e, m = u27_fields(u27)
671     if e == 0:
672         return -0.0 if s else 0.0
673     return (-1.0 if s else 1.0) * (2.0 ** (int(e) - BIAS)) * (1.0 + (m / (1 << M18_BITS)))
674
675 def u27_from_float_trunc(x: float) -> int:
676     """TRUNC into S1E8M18. Denorm flush to 0. Overflow clamp."""
677     if x == 0.0 or not np.isfinite(x):
678         if x == 0.0:
679             return 0
680         sign = 1 if np.signbit(x) else 0
681         return u27_pack(sign, 254, M18_MASK)
682
683     sign = 1 if x < 0 else 0
684     ax = abs(x)
685     if ax == 0.0:
686         return 0
687
688     m, e = np.frexp(ax)
689     m2 = m * 2.0
690     e2 = e - 1
691     exp = int(e2 + BIAS)
692
693     if exp <= 0:
694         return 0
695     if exp >= 255:
696         return u27_pack(sign, 254, M18_MASK)
697
698     mant = int(np.floor((m2 - 1.0) * (1 << M18_BITS)))
699     if mant < 0:
700         mant = 0
701     if mant > M18_MASK:
702         mant = M18_MASK
703
704     return u27_pack(sign, exp, mant)
705
706 # =====
707 # File IO helpers
708 # =====
709 def load_frame0_u27(path: str):
710     idx, px, py, vx, vy, mm = [], [], [], [], [], []
711     with open(path, "r") as f:
712         for line in f:
713             line = line.strip()
714             if (not line) or line.startswith("#"):
715                 continue
716             parts = line.split()
717             idx.append(int(parts[0]))
718             px.append(u27_from_hex(parts[1]))
719             py.append(u27_from_hex(parts[2]))

```

```

720         vx.append(u27_from_hex(parts[3]))
721         vy.append(u27_from_hex(parts[4]))
722         mm.append(u27_from_hex(parts[5]))
723     return (
724         np.array(idx, dtype=np.int32),
725         np.array(px, dtype=np.uint32),
726         np.array(py, dtype=np.uint32),
727         np.array(vx, dtype=np.uint32),
728         np.array(vy, dtype=np.uint32),
729         np.array(mm, dtype=np.uint32),
730     )
731
732 def write_frame_chip_hex(path: str, idx, px_u27, py_u27, vx_u27, vy_u27, m_u27):
733     with open(path, "w") as f:
734         f.write("# i px py vx vy m (S1E8M18 hex)\n")
735         for i in range(len(idx)):
736             f.write(
737                 f"{int(idx[i])} {u27_to_hex(int(px_u27[i]))} {u27_to_hex(int(py_u27[i]))} "
738                 f"{u27_to_hex(int(vx_u27[i]))} {u27_to_hex(int(vy_u27[i]))} {u27_to_hex("
739                 f"int(m_u27[i]))}\n"
740             )
741
742 def write_accel_core27_hex(path: str, ax_u27, ay_u27):
743     with open(path, "w") as f:
744         f.write("# i ax27 ay27 (S1E8M18 hex)\n")
745         for i in range(len(ax_u27)):
746             f.write(f"{i} {u27_to_hex(int(ax_u27[i]))} {u27_to_hex(int(ay_u27[i]))}\n")
747
748 def write_frame_host_decimal(path: str, idx, px_f, py_f, vx_f, vy_f, m_f, g0: float, dt:
749 float, k_host: float):
750     with open(path, "w") as f:
751         f.write(f"# G0 = {g0:.12e}\n")
752         f.write(f"# DT = {dt:.12e}\n")
753         f.write(f"# K_HOST = G0*DT = {k_host:.12e}\n")
754         f.write("# i px py vx vy m (decimal float)\n")
755         for i in range(len(idx)):
756             f.write(
757                 f"{int(idx[i])} {float(px_f[i]): .12e} {float(py_f[i]): .12e} "
758                 f"{float(vx_f[i]): .12e} {float(vy_f[i]): .12e} {float(m_f[i]): .12e}\n"
759             )
760
761 # =====
762 # Core compute: chip hex -> internal u27 -> accel
763 # =====
764 def compute_acc_once_u27io_accum27(px_u27, py_u27, m_u27, eps2_u27: int):
765     """
766     - input is already 27-bit S1E8M18
767     - all ops in 27-bit bit-exact core
768     - sum_ax/sum_ay accumulated in 27-bit
769     """
770     N = len(px_u27)
771     ax_u27 = np.zeros(N, dtype=np.uint32)
772     ay_u27 = np.zeros(N, dtype=np.uint32)
773
774     px = [int(v) & U27_MASK for v in px_u27]
775     py = [int(v) & U27_MASK for v in py_u27]
776     mm = [int(v) & U27_MASK for v in m_u27]
777
778     for i in range(N):
779         sum_ax = 0
780         sum_ay = 0
781
782         for j in range(N):
783             if j == i:
784                 continue
785
786             dx = fp27_add(px[j], u27_neg(px[i]))
787             dy = fp27_add(py[j], u27_neg(py[i]))
788
789             dx2 = fp27_mul(dx, dx)
790             dy2 = fp27_mul(dy, dy)
791             r2 = fp27_add(dx2, dy2)

```

```

790         r2e = fp27_add(r2, eps2_u27)
791
792         # FIX: match RTL two_body_core.v pipeline mul order exactly:
793         #   c12: s2 = s*s AND t = m2*s (parallel)
794         #   c13: k = t*s2 (m2*s) * (s*s)
795         # NOT the same as m2 * (s*s*s) due to FpMul bit0 truncation
796         s = fast_inv_sqrt_u27(r2e)
797         s2 = fp27_mul(s, s) # s
798         t = fp27_mul(mm[j], s) # m s (parallel with s2 in RTL)
799         k = fp27_mul(t, s2) # ( m s ) s
800
801         term_x = fp27_mul(k, dx)
802         term_y = fp27_mul(k, dy)
803
804         sum_ax = fp27_add(sum_ax, term_x)
805         sum_ay = fp27_add(sum_ay, term_y)
806
807         ax_u27[i] = sum_ax & U27_MASK
808         ay_u27[i] = sum_ay & U27_MASK
809
810     return ax_u27, ay_u27
811
812     # =====
813     # Main loop
814     # =====
815     def main():
816         os.makedirs(DIR_FRAME_HOST, exist_ok=True)
817         os.makedirs(DIR_FRAME_CHIP, exist_ok=True)
818         os.makedirs(DIR_ACCEL_CHIP, exist_ok=True)
819         os.makedirs(DIR_ACCEL_CORE27, exist_ok=True)
820
821         print("Imported core modules:")
822         print(" fp27_add ->", fp27_add.__module__)
823         print(" fp27_mul ->", fp27_mul.__module__)
824         print(" fast_inv_sqrt_u27 ->", fast_inv_sqrt_u27.__module__)
825
826         eps2_u27 = EPS_SQUARE_U27
827
828         # Load frame0 (27-bit hex)
829         idx, px0_u27, py0_u27, vx0_u27, vy0_u27, m0_u27 = load_frame0_u27(INPUT_FILE)
830
831         # Copy frame0 into frame_chip_input/frame0.txt
832         frame0_chip_path = os.path.join(DIR_FRAME_CHIP, "frame0.txt")
833         try:
834             shutil.copyfile(INPUT_FILE, frame0_chip_path)
835         except Exception:
836             write_frame_chip_hex(frame0_chip_path, idx, px0_u27, py0_u27, vx0_u27, vy0_u27,
837                                 m0_u27)
838
839         # Host state (float) initialized from frame0 hex --- FP32
840         px_f = np.array([u27_to_float(int(v)) for v in px0_u27], dtype=HOST_DTYPE)
841         py_f = np.array([u27_to_float(int(v)) for v in py0_u27], dtype=HOST_DTYPE)
842         vx_f = np.array([u27_to_float(int(v)) for v in vx0_u27], dtype=HOST_DTYPE)
843         vy_f = np.array([u27_to_float(int(v)) for v in vy0_u27], dtype=HOST_DTYPE)
844         m_f = np.array([u27_to_float(int(v)) for v in m0_u27], dtype=HOST_DTYPE)
845
846         # Chip state (hex) initialized from frame0
847         px_u27 = px0_u27.copy()
848         py_u27 = py0_u27.copy()
849         vx_u27 = vx0_u27.copy()
850         vy_u27 = vy0_u27.copy()
851         m_u27 = m0_u27.copy()
852
853         # Dump host decimal frame0
854         host0_path = os.path.join(DIR_FRAME_HOST, "frame0.txt")
855         write_frame_host_decimal(host0_path, idx, px_f, py_f, vx_f, vy_f, m_f, GO, DT, K_HOST)
856
857         # Produce frames 0..N_FRAMES, and accels 0..N_FRAMES-1
858         for k in range(N_FRAMES):
859             # 1) Chip computes accel for current frame k (NO GO inside)
860             ax_u27, ay_u27 = compute_acc_once_u27io_accum27(
861                 px_u27, py_u27, m_u27, eps2_u27
862             )

```

```

862
863     accel_path = os.path.join(DIR_ACCEL_CHIP, f"accel{k}.txt")
864     write_accel_core27_hex(accel_path, ax_u27, ay_u27)
865
866     accel27_path = os.path.join(DIR_ACCEL_CORE27, f"accel27_{k}.txt")
867     write_accel_core27_hex(accel27_path, ax_u27, ay_u27)
868
869     # 2) Host update (FP32) using accel from chip (apply K_HOST only here)
870     ax_f = np.array([u27_to_float(int(v)) for v in ax_u27], dtype=HOST_DTYPE)
871     ay_f = np.array([u27_to_float(int(v)) for v in ay_u27], dtype=HOST_DTYPE)
872
873     k_host_f = HOST_DTYPE(K_HOST)
874     dt_f      = HOST_DTYPE(DT)
875
876     vx_f = (vx_f + ax_f * k_host_f).astype(HOST_DTYPE, copy=False)
877     vy_f = (vy_f + ay_f * k_host_f).astype(HOST_DTYPE, copy=False)
878     px_f = (px_f + vx_f * dt_f).astype(HOST_DTYPE, copy=False)
879     py_f = (py_f + vy_f * dt_f).astype(HOST_DTYPE, copy=False)
880
881     # 3) Quantize updated host frame back to 27-bit chip hex for next frame using TRUNC
882     px_u27 = np.array([u27_from_float_trunc(float(x)) for x in px_f], dtype=np.uint32)
883     py_u27 = np.array([u27_from_float_trunc(float(x)) for x in py_f], dtype=np.uint32)
884     vx_u27 = np.array([u27_from_float_trunc(float(x)) for x in vx_f], dtype=np.uint32)
885     vy_u27 = np.array([u27_from_float_trunc(float(x)) for x in vy_f], dtype=np.uint32)
886
887     # mass stays constant
888     m_u27 = m_u27
889
890     # 4) Write next chip input frame and host decimal frame
891     chip_path = os.path.join(DIR_FRAME_CHIP, f"frame{k+1}.txt")
892     write_frame_chip_hex(chip_path, idx, px_u27, py_u27, vx_u27, vy_u27, m_u27)
893
894     host_path = os.path.join(DIR_FRAME_HOST, f"frame{k+1}.txt")
895     write_frame_host_decimal(host_path, idx, px_f, py_f, vx_f, vy_f, m_f, GO, DT, K_HOST
)
896
897     print(f"[OK] Done. OUT_ROOT={OUT_ROOT}")
898     print(f"     EPS_SQUARE_U27=0x{eps2_u27:07X}")
899     print(f"     GO={GO} DT={DT} K_HOST={K_HOST}")
900     print(f"     HOST_DTYPE={HOST_DTYPE}")
901     print(f"     Frames written: frame0..frame{N_FRAMES}")
902     print(f"     Accels written: accel0..accel{N_FRAMES-1}")
903     print(f"     Core27 written: accel27_0..accel27_{N_FRAMES-1}")
904
905     # if __name__ == "__main__":
906     #     main()
907
908
909     # =====
910     # TB-compatible one-shot golden output
911     # Match tb_core_accel input/output behavior
912     # =====
913
914     TB_INPUT_FILE = "../Hardware/tb/frame_input/frame0_1024binit200_27bits.txt"
915     TB_GOLDEN_OUT = "../Hardware/tb/output/golden_accel_1024binit200_27bits.txt"
916
917     def run_tb_compatible_golden():
918         idx, px_u27, py_u27, vx_u27, vy_u27, m_u27 = load_frame0_u27(TB_INPUT_FILE)
919
920         ax_u27, ay_u27 = compute_acc_once_u27io_accum27(
921             px_u27,
922             py_u27,
923             m_u27,
924             EPS_SQUARE_U27
925         )
926
927         with open(TB_GOLDEN_OUT, "w") as f:
928             f.write("# i ax ay (S1E8M18 hex)\n")
929             for i in range(len(idx)):
930                 f.write(f"{int(idx[i]):4d}  {u27_to_hex(int(ax_u27[i]))}  {u27_to_hex(int(ay_u27
931 [i]))}\n")
932
933     print(f"[OK] Wrote TB-compatible golden output: {TB_GOLDEN_OUT}")

```

```

933
934 if __name__ == "__main__":
935     run_tb_compatible_golden()
936
937
938 # ## 5. Compare
939 import os
940
941 # =====
942 # USER CONFIG
943 # =====
944 RTL_FILE = "../Hardware/tb/output/eps025_accel_1024binit200_27bits.txt"
945 PY_FILE = "../Hardware/tb/output/golden_accel_1024binit200_27bits.txt"
946 # RTL_FILE = "eps025_temp_27bits_8binit10.txt"
947 # PY_FILE = "eps025_v4real_chiplate_8binit10/accel_core27_output/accel27_0.txt"
948
949 # if you want, save mismatch report
950 REPORT_FILE = "compare_27bit_report.txt"
951
952 # =====
953 # HELPERS
954 # =====
955 U27_MASK = (1 << 27) - 1
956
957 def popcount(x: int) -> int:
958     return bin(x).count("1")
959
960 def load_27bit_txt(path: str):
961     """
962     Expected format:
963     # i ax27 ay27 ...
964     0 3ABCDE1 0123456
965     1 1234567 7654321
966     Returns:
967     dict[idx] = (ax27_int, ay27_int)
968     """
969     data = {}
970     with open(path, "r") as f:
971         for line in f:
972             line = line.strip()
973             if not line or line.startswith("#"):
974                 continue
975             parts = line.split()
976             if len(parts) < 3:
977                 continue
978             idx = int(parts[0])
979             ax = int(parts[1], 16) & U27_MASK
980             ay = int(parts[2], 16) & U27_MASK
981             data[idx] = (ax, ay)
982     return data
983
984 # =====
985 # MAIN
986 # =====
987 def main():
988     rtl = load_27bit_txt(RTL_FILE)
989     py = load_27bit_txt(PY_FILE)
990
991     rtl_keys = set(rtl.keys())
992     py_keys = set(py.keys())
993
994     only_in_rtl = sorted(rtl_keys - py_keys)
995     only_in_py = sorted(py_keys - rtl_keys)
996     common = sorted(rtl_keys & py_keys)
997
998     lines = []
999     lines.append(f"RTL_FILE = {RTL_FILE}")
1000    lines.append(f"PY_FILE = {PY_FILE}")
1001    lines.append("")
1002
1003    if only_in_rtl:
1004        lines.append(f"indices only in RTL: {only_in_rtl}")
1005    if only_in_py:

```

```

1006     lines.append(f"indices only in PY : {only_in_py}")
1007     if only_in_rtl or only_in_py:
1008         lines.append("")
1009
1010     mismatch_count = 0
1011     max_ax_bits = -1
1012     max_ay_bits = -1
1013     max_ax_idx = None
1014     max_ay_idx = None
1015
1016     total_ax_bits = 0
1017     total_ay_bits = 0
1018
1019     for i in common:
1020         rtl_ax, rtl_ay = rtl[i]
1021         py_ax, py_ay = py[i]
1022
1023         diff_ax = rtl_ax ^ py_ax
1024         diff_ay = rtl_ay ^ py_ay
1025
1026         bits_ax = popcount(diff_ax)
1027         bits_ay = popcount(diff_ay)
1028
1029         total_ax_bits += bits_ax
1030         total_ay_bits += bits_ay
1031
1032         if bits_ax > max_ax_bits:
1033             max_ax_bits = bits_ax
1034             max_ax_idx = i
1035
1036         if bits_ay > max_ay_bits:
1037             max_ay_bits = bits_ay
1038             max_ay_idx = i
1039
1040         if diff_ax != 0 or diff_ay != 0:
1041             mismatch_count += 1
1042             lines.append(
1043                 f"{i:4d} | "
1044                 f"{rtl_ax:07X},{rtl_ay:07X} -> {py_ax:07X},{py_ay:07X} | "
1045                 f"{bits_ax},{bits_ay}"
1046             )
1047
1048     lines.append("")
1049     lines.append("Summary")
1050     lines.append("-----")
1051     lines.append(f"common indices      : {len(common)}")
1052     lines.append(f"mismatch rows      : {mismatch_count}")
1053     lines.append(f"total ax bit diffs : {total_ax_bits}")
1054     lines.append(f"total ay bit diffs : {total_ay_bits}")
1055     lines.append(f"max ax bit diff    : {max_ax_bits} (index {max_ax_idx})")
1056     lines.append(f"max ay bit diff    : {max_ay_bits} (index {max_ay_idx})")
1057
1058     report = "\n".join(lines)
1059     print(report)
1060
1061     with open(REPORT_FILE, "w") as f:
1062         f.write(report)
1063
1064     print(f"\nWrote report to: {REPORT_FILE}")
1065
1066 # if __name__ == "__main__":
1067 #     main()

```

## A.5.2 golden\_control.py

Listing 42: code/Golden/golden\_control.py

```

1  #!/usr/bin/env python3
2  """

```

```

3 Golden model for nbody_control + four_core_wrapper + nbody_integrator.
4
5 Default output matches Hardware/tb/tb_nbody_control.sv:
6     # i x y vx vy ax ay (S1E8M18 hex)
7         0 0123456 0ABCDEF 0000000 0000000 0123000 0ABC000
8
9 The model reuses the bit-exact FP/core helpers from golden.py, then applies the
10 same control-level ordering: 16-body tiles, four 4-lane groups per tile, and
11 ascending lane writeback. It also updates the body memory with the RTL
12 integrator semantics so --gap and --first-step match nbody_control.
13 """
14
15 import argparse
16 import os
17 import sys
18 from typing import Iterable, List, Optional, Tuple
19
20
21 SCRIPT_DIR = os.path.dirname(os.path.abspath(__file__))
22 REPO_ROOT = os.path.abspath(os.path.join(SCRIPT_DIR, ".."))
23 if SCRIPT_DIR not in sys.path:
24     sys.path.insert(0, SCRIPT_DIR)
25
26 import golden as core # noqa: E402
27
28
29 U27_MASK = (1 << 27) - 1
30 SIGN_SHIFT = 26
31 EXP_SHIFT = 18
32 E_MASK = 0xFF
33 M18_MASK = (1 << 18) - 1
34 TILE_STRIDE = 16
35 GROUP_SIZE = 4
36 EPS_SQUARE_U27 = (0 << 26) | (125 << 18) | 0
37
38 DEFAULT_INPUT = os.path.join(
39     REPO_ROOT, "Hardware", "tb", "frame_input", "frame0_1024binit200_27bits.txt"
40 )
41 DEFAULT_STATE_OUT = os.path.join(
42     REPO_ROOT, "Hardware", "tb", "output", "golden_control_state_1024binit200_27bits.txt"
43 )
44
45
46 class BodyState:
47     def __init__(
48         self,
49         x: List[int],
50         y: List[int],
51         vx: List[int],
52         vy: List[int],
53         m: List[int],
54         ax: List[int],
55         ay: List[int],
56     ) -> None:
57         self.x = x
58         self.y = y
59         self.vx = vx
60         self.vy = vy
61         self.m = m
62         self.ax = ax
63         self.ay = ay
64
65
66 def u27_hex(value: int) -> str:
67     return f"{value & U27_MASK:07X}"
68
69
70 def u27_from_hex(text: str) -> int:
71     return int(text, 16) & U27_MASK
72
73
74 def u27_pack(sign: int, exp: int, mant: int) -> int:
75     return ((sign & 1) << SIGN_SHIFT) | ((exp & E_MASK) << EXP_SHIFT) | (mant & M18_MASK)

```

```

76
77
78 def u27_fields(value: int) -> Tuple[int, int, int]:
79     value &= U27_MASK
80     return (value >> SIGN_SHIFT) & 1, (value >> EXP_SHIFT) & E_MASK, value & M18_MASK
81
82
83 def half_step_accel(value: int) -> int:
84     """Match nbody_control's exponent-only half-step acceleration adjustment."""
85     sign, exp, mant = u27_fields(value)
86     half_exp = exp - 1 if exp > 1 else 0
87     return u27_pack(sign, half_exp, mant)
88
89
90 def integrator_step(
91     x: int,
92     y: int,
93     vx: int,
94     vy: int,
95     ax: int,
96     ay: int,
97     half_step: bool = False,
98 ) -> Tuple[int, int, int, int]:
99     """Match nbody_integrator.sv: vx'=vx+ax, x'=x+vx'."""
100     if half_step:
101         ax = half_step_accel(ax)
102         ay = half_step_accel(ay)
103
104     vx_new = core.fp27_add_rtl(vx, ax)
105     vy_new = core.fp27_add_rtl(vy, ay)
106     x_new = core.fp27_add_rtl(x, vx_new)
107     y_new = core.fp27_add_rtl(y, vy_new)
108     return x_new, y_new, vx_new, vy_new
109
110
111 def load_frame(path: str, max_bodies: int) -> Tuple[BodyState, int]:
112     x = [0] * max_bodies
113     y = [0] * max_bodies
114     vx = [0] * max_bodies
115     vy = [0] * max_bodies
116     m = [0] * max_bodies
117     highest_idx = -1
118
119     with open(path, "r") as f:
120         for line_no, line in enumerate(f, start=1):
121             line = line.strip()
122             if not line or line.startswith("#"):
123                 continue
124
125             parts = line.split()
126             if len(parts) < 6:
127                 raise ValueError(f"{path}:{line_no}: expected 'idx x y vx vy m'")
128
129             idx = int(parts[0])
130             if idx < 0 or idx >= max_bodies:
131                 continue
132
133             x[idx] = u27_from_hex(parts[1])
134             y[idx] = u27_from_hex(parts[2])
135             vx[idx] = u27_from_hex(parts[3])
136             vy[idx] = u27_from_hex(parts[4])
137             m[idx] = u27_from_hex(parts[5])
138             highest_idx = max(highest_idx, idx)
139
140     return BodyState(x=x, y=y, vx=vx, vy=vy, m=m, ax=[0] * max_bodies, ay=[0] * max_bodies),
141         highest_idx + 1
142
143 def compute_one_body_accel(state: BodyState, i: int, active_count: int) -> Tuple[int, int]:
144     """Bit-exact model of one accumulated four_core_wrapper output lane."""
145     sum_ax = 0
146     sum_ay = 0
147

```

```

148     for j in range(active_count):
149         if j == i:
150             continue
151
152         dx = core.fp27_add_rtl(state.x[j], core.u27_neg(state.x[i]))
153         dy = core.fp27_add_rtl(state.y[j], core.u27_neg(state.y[i]))
154
155         dx2 = core.fp27_mul_rtl(dx, dx)
156         dy2 = core.fp27_mul_rtl(dy, dy)
157         r2 = core.fp27_add_rtl(dx2, dy2)
158         r2e = core.fp27_add_rtl(r2, EPS_SQUARE_U27)
159
160         s = core.fast_inv_sqrt_u27_rtl(r2e)
161         s2 = core.fp27_mul_rtl(s, s)
162         t = core.fp27_mul_rtl(state.m[j], s)
163         k = core.fp27_mul_rtl(t, s2)
164
165         term_x = core.fp27_mul_rtl(k, dx)
166         term_y = core.fp27_mul_rtl(k, dy)
167         sum_ax = core.fp27_add_rtl(sum_ax, term_x)
168         sum_ay = core.fp27_add_rtl(sum_ay, term_y)
169
170     return sum_ax & U27_MASK, sum_ay & U27_MASK
171
172
173 def control_write_order(active_count: int) -> Iterable[int]:
174     """Yield body indices in the same tile/group/lane order as nbody_control."""
175     for tile_base in range(0, active_count, TILE_STRIDE):
176         for group in range(TILE_STRIDE // GROUP_SIZE):
177             group_base = tile_base + group * GROUP_SIZE
178             for lane in range(GROUP_SIZE):
179                 idx = group_base + lane
180                 if idx < active_count:
181                     yield idx
182
183
184 def compute_accel_pass(state: BodyState, active_count: int) -> List[Tuple[int, int, int]]:
185     writes: List[Tuple[int, int, int]] = []
186
187     for i in control_write_order(active_count):
188         ax, ay = compute_one_body_accel(state, i, active_count)
189         state.ax[i] = ax
190         state.ay[i] = ay
191         writes.append((i, ax, ay))
192
193     return writes
194
195
196 def integrate_pass(state: BodyState, active_count: int, half_step: bool) -> None:
197     for i in range(active_count):
198         x, y, vx, vy = integrator_step(
199             state.x[i],
200             state.y[i],
201             state.vx[i],
202             state.vy[i],
203             state.ax[i],
204             state.ay[i],
205             half_step=half_step,
206         )
207         state.x[i] = x
208         state.y[i] = y
209         state.vx[i] = vx
210         state.vy[i] = vy
211
212
213 def run_control_model(
214     state: BodyState,
215     active_count: int,
216     gap: int,
217     first_step: bool,
218 ) -> Tuple[List[Tuple[int, int, int]], BodyState]:
219     """
220     Run the control-level model.

```

```

221
222 Returns the first timestep's accel writes because tb_nbody_control.sv uses
223 gap=1 and writes one flat accel output file. For gap>1, later timesteps are
224 still integrated into state, but not appended to that flat file.
225 """
226 if gap <= 0 or active_count <= 0:
227     return [], state
228
229 first_writes: Optional[List[Tuple[int, int, int]]] = None
230 for timestep in range(gap):
231     writes = compute_accel_pass(state, active_count)
232     if first_writes is None:
233         first_writes = writes
234     integrate_pass(state, active_count, half_step=first_step and timestep == 0)
235
236 return first_writes or [], state
237
238
239 def write_accel_file(path: str, writes: Iterable[Tuple[int, int, int]]) -> None:
240     os.makedirs(os.path.dirname(os.path.abspath(path)), exist_ok=True)
241     with open(path, "w") as f:
242         f.write("# i ax ay (S1E8M18 hex)\n")
243         for idx, ax, ay in writes:
244             f.write(f"{idx:4d} {u27_hex(ax)} {u27_hex(ay)}\n")
245
246
247 def write_frame_file(path: str, state: BodyState, active_count: int) -> None:
248     os.makedirs(os.path.dirname(os.path.abspath(path)), exist_ok=True)
249     with open(path, "w") as f:
250         f.write("# i px py vx vy m (S1E8M18 hex)\n")
251         for i in range(active_count):
252             f.write(
253                 f"{i:4d} {u27_hex(state.x[i])} {u27_hex(state.y[i])} "
254                 f"{u27_hex(state.vx[i])} {u27_hex(state.vy[i])} {u27_hex(state.m[i])}\n"
255             )
256
257
258 def write_state_file(path: str, state: BodyState, active_count: int) -> None:
259     os.makedirs(os.path.dirname(os.path.abspath(path)), exist_ok=True)
260     with open(path, "w") as f:
261         f.write("# i x y vx vy ax ay (S1E8M18 hex)\n")
262         for i in range(active_count):
263             f.write(
264                 f"{i:4d} {u27_hex(state.x[i])} {u27_hex(state.y[i])} "
265                 f"{u27_hex(state.vx[i])} {u27_hex(state.vy[i])} "
266                 f"{u27_hex(state.ax[i])} {u27_hex(state.ay[i])}\n"
267             )
268
269
270 def parse_args() -> argparse.Namespace:
271     parser = argparse.ArgumentParser(
272         description="Generate a golden output for nbody_control.sv."
273     )
274     parser.add_argument("--input", default=DEFAULT_INPUT, help="Input frame file.")
275     parser.add_argument(
276         "--output",
277         default=DEFAULT_STATE_OUT,
278         help="Output final x/y/vx/vy/ax/ay file matching tb_nbody_control.sv.",
279     )
280     parser.add_argument("--frame-output", default=None, help="Optional final updated frame output.")
281     parser.add_argument(
282         "--accel-output",
283         default=None,
284         help="Optional acceleration-only write stream file.",
285     )
286     parser.add_argument("--max-bodies", type=int, default=1024, help="Memory capacity.")
287     parser.add_argument("--n-bodies", type=int, default=None, help="Active body count.")
288     parser.add_argument("--gap", type=int, default=1, help="Number of timesteps to run.")
289     parser.add_argument(
290         "--first-step",
291         action="store_true",
292         help="Apply the initial leapfrog half-step kick on timestep 0.",

```

```

293 )
294 return parser.parse_args()
295
296
297 def main() -> None:
298     args = parse_args()
299     state, loaded_count = load_frame(args.input, args.max_bodies)
300     active_count = args.n_bodies if args.n_bodies is not None else loaded_count
301     active_count = max(0, min(active_count, args.max_bodies))
302
303     writes, final_state = run_control_model(
304         state=state,
305         active_count=active_count,
306         gap=args.gap,
307         first_step=args.first_step,
308     )
309     write_state_file(args.output, final_state, active_count)
310
311     if args.frame_output:
312         write_frame_file(args.frame_output, final_state, active_count)
313
314     if args.accel_output:
315         write_accel_file(args.accel_output, writes)
316
317     print(f"[OK] input      : {args.input}")
318     print(f"[OK] n_bodies    : {active_count}")
319     print(f"[OK] gap          : {args.gap}")
320     print(f"[OK] first_step   : {int(args.first_step)}")
321     print(f"[OK] state output  : {args.output}")
322     if args.frame_output:
323         print(f"[OK] frame output : {args.frame_output}")
324     if args.accel_output:
325         print(f"[OK] accel output : {args.accel_output}")
326
327
328 if __name__ == "__main__":
329     main()

```

### A.5.3 golden\_avmm\_xy.py

Listing 43: code/Golden/golden\_avmm\_xy.py

```

1  #!/usr/bin/env python3
2  """
3  Golden X/Y stream for the current nbody_accel_avmm hardware path.
4
5  This script takes the same six-column S1E8M18 frame input file that software
6  loads into the accelerator, runs the bit-level Python model of the current RTL,
7  and writes the X/Y positions that software reads from OUT_X/OUT_Y after DONE.
8
9  The default first-step behavior matches nbody_accel_avmm.sv: loading bodies via
10 VY_IN marks the next GO as the initial leapfrog half-step run.
11 """
12
13 import argparse
14 import os
15 import sys
16 from typing import Optional
17
18
19 SCRIPT_DIR = os.path.dirname(os.path.abspath(__file__))
20 REPO_ROOT = os.path.abspath(os.path.join(SCRIPT_DIR, ".."))
21 if SCRIPT_DIR not in sys.path:
22     sys.path.insert(0, SCRIPT_DIR)
23
24 import golden_control as hw # noqa: E402
25
26
27 DEFAULT_INPUT = os.path.join(

```

```

28     REPO_ROOT, "Hardware", "tb", "frame_input", "frame0_1024binit200_27bits.txt"
29 )
30 DEFAULT_OUTPUT = os.path.join(
31     REPO_ROOT, "Hardware", "tb", "output", "golden_xy_1024binit200_27bits.txt"
32 )
33
34
35 def checked_active_count(loaded_count: int, requested_count: Optional[int], max_bodies: int)
36     -> int:
37     if requested_count is None:
38         active_count = loaded_count
39     else:
40         active_count = requested_count
41
42     if active_count <= 0:
43         raise ValueError(f"n_bodies must be positive, got {active_count}")
44     if active_count > max_bodies:
45         raise ValueError(f"n_bodies={active_count} exceeds max_bodies={max_bodies}")
46     if active_count > loaded_count:
47         raise ValueError(
48             f"n_bodies={active_count} but input only contains rows through index {
49             loaded_count - 1}"
50         )
51
52     return active_count
53
54 def write_xy_file(path: str, input_path: str, active_count: int, gap: int, first_step: bool,
55     state: hw.BodyState) -> None:
56     out_dir = os.path.dirname(os.path.abspath(path))
57     if out_dir:
58         os.makedirs(out_dir, exist_ok=True)
59
60     with open(path, "w") as f:
61         f.write(f"# source {input_path}\n")
62         f.write(f"# n_bodies {active_count}\n")
63         f.write(f"# gap {gap}\n")
64         f.write(f"# first_step {int(first_step)}\n")
65         f.write("# i x y (S1E8M18 hex)\n")
66         for i in range(active_count):
67             f.write(f"{i:4d} {hw.u27_hex(state.x[i])} {hw.u27_hex(state.y[i])}\n")
68
69 def parse_args() -> argparse.Namespace:
70     parser = argparse.ArgumentParser(
71         description=(
72             "Generate the golden OUT_X/OUT_Y stream returned by current "
73             "nbody_accel_avmm hardware after one GO/DONE run."
74         )
75     )
76     parser.add_argument("--input", default=DEFAULT_INPUT, help="Input frame txt file.")
77     parser.add_argument("--output", default=DEFAULT_OUTPUT, help="Output X/Y txt file.")
78     parser.add_argument("--max-bodies", type=int, default=1024, help="Hardware memory
79     capacity.")
80     parser.add_argument("--n-bodies", type=int, default=None, help="Active body count.")
81     parser.add_argument("--gap", type=int, default=1, help="Hardware GAP register value.")
82     parser.add_argument(
83         "--no-first-step",
84         action="store_true",
85         help="Model a later GO without reloading bodies. Default models GO after body load."
86     )
87     parser.add_argument(
88         "--frame-output",
89         default=None,
90         help="Optional full final frame output with x/y/vx/vy/m.",
91     )
92     parser.add_argument(
93         "--accel-output",
94         default=None,
95         help="Optional acceleration write stream output, matching tb_nbody_control.sv.",
96     )
97     return parser.parse_args()

```

```

97
98
99 def main() -> None:
100     args = parse_args()
101
102     if args.max_bodies <= 0:
103         raise ValueError(f"max_bodies must be positive, got {args.max_bodies}")
104     if args.gap <= 0:
105         raise ValueError(f"gap must be positive, got {args.gap}")
106
107     state, loaded_count = hw.load_frame(args.input, args.max_bodies)
108     active_count = checked_active_count(loaded_count, args.n_bodies, args.max_bodies)
109     first_step = not args.no_first_step
110
111     accel_writes, final_state = hw.run_control_model(
112         state=state,
113         active_count=active_count,
114         gap=args.gap,
115         first_step=first_step,
116     )
117
118     write_xy_file(
119         path=args.output,
120         input_path=args.input,
121         active_count=active_count,
122         gap=args.gap,
123         first_step=first_step,
124         state=final_state,
125     )
126
127     if args.frame_output:
128         hw.write_frame_file(args.frame_output, final_state, active_count)
129
130     if args.accel_output:
131         hw.write_accel_file(args.accel_output, accel_writes)
132
133     print(f"[OK] input      : {args.input}")
134     print(f"[OK] n_bodies    : {active_count}")
135     print(f"[OK] gap          : {args.gap}")
136     print(f"[OK] first_step   : {int(first_step)}")
137     print(f"[OK] xy output    : {args.output}")
138     if args.frame_output:
139         print(f"[OK] frame output : {args.frame_output}")
140     if args.accel_output:
141         print(f"[OK] accel output : {args.accel_output}")
142
143
144 if __name__ == "__main__":
145     main()

```

#### A.5.4 FP64.py

Listing 44: code/Golden/FP64.py

```

1  # ## 2. FP64 Golden Reference
2
3
4  # fp64_reference_loop.py
5  # Pure float64 N-body reference with "chip-accel convention":
6  #   - accel_chip_output: DOES NOT include GO
7  #   - host update uses:  K_HOST = GO * DT   (v += a_chip * K_HOST)
8  #
9  # Output structure same as chiplike version (3 folders)
10
11 import os
12 import shutil
13 import numpy as np
14
15 # =====

```

```

16 # User knobs
17 # =====
18 INPUT_FILE = "../Hardware/tb/frame_input/frame0_256binit200.txt" # f187 hex input
19
20 OUT_ROOT = "../Hardware/tb/output/tempfp64_256binit200_outputs"
21 DIR_FRAME_HOST = os.path.join(OUT_ROOT, "frame_host_output")
22 DIR_FRAME_CHIP = os.path.join(OUT_ROOT, "frame_chip_input")
23 DIR_ACCEL_CHIP = os.path.join(OUT_ROOT, "accel_chip_output")
24
25 N_FRAMES = 10
26 DT = 0.1
27 GO = 10.0
28
29 # IMPORTANT: this is epsilon^2 added to r^2, use very small value for closer-to-realistic
    accel (but not too small to cause underflow issues in fp64)
30 EPS_SQ = 2.0 ** (-27)
31
32 # Host update gain (this is where GO goes)
33 K_HOST = GO * DT
34
35 # =====
36 # SiE8M7 helpers (same as before)
37 # =====
38 BIAS = 127
39 E_BITS = 8
40 M_BITS = 7
41 TOTAL_BITS = 16
42 E_MASK = (1 << E_BITS) - 1
43 M_MASK = (1 << M_BITS) - 1
44 SIGN_SHIFT = TOTAL_BITS - 1
45 EXP_SHIFT = M_BITS
46 U16_MASK = (1 << TOTAL_BITS) - 1
47
48 def f16_pack(sign: int, exp: int, mant: int) -> int:
49     return ((sign & 1) << SIGN_SHIFT) | ((exp & E_MASK) << EXP_SHIFT) | (mant & M_MASK)
50
51 def f16_to_hex(u16: int) -> str:
52     return f"{u16 & U16_MASK:04X}"
53
54 def f16_from_hex(h: str) -> int:
55     return int(h, 16) & U16_MASK
56
57 def f16_to_float(u16: int) -> float:
58     if u16 == 0:
59         return 0.0
60     sign = (u16 >> SIGN_SHIFT) & 1
61     exp = (u16 >> EXP_SHIFT) & E_MASK
62     mant = u16 & M_MASK
63     if exp == 0:
64         return 0.0
65     val = (1.0 + mant / (1 << M_BITS)) * (2.0 ** (exp - BIAS))
66     return -val if sign else val
67
68 def f16_from_float(x: float) -> int:
69     if x == 0.0:
70         return 0
71     sign = 1 if x < 0 else 0
72     ax = abs(x)
73
74     m, e = np.frexp(ax) # ax = m * 2^e, m in [0.5, 1)
75     m2 = m * 2.0
76     e2 = e - 1
77     exp = e2 + BIAS
78
79     if exp <= 0:
80         return 0
81     if exp >= 255:
82         return f16_pack(sign, 254, M_MASK)
83
84     frac = m2 - 1.0
85     mant_f = frac * (1 << M_BITS)
86     mant = int(np.floor(mant_f + 0.5)) # round-to-nearest
87

```

```

88     if mant == (1 << M_BITS):
89         mant = 0
90         exp += 1
91         if exp >= 255:
92             exp = 254
93             mant = M_MASK
94
95     return f16_pack(sign, exp, mant)
96
97 # =====
98 # IO helpers
99 # =====
100 def load_frame0_hex(path: str):
101     idx, px, py, vx, vy, mm = [], [], [], [], [], []
102     with open(path, "r") as f:
103         for line in f:
104             if line.startswith("#") or not line.strip():
105                 continue
106             parts = line.split()
107             idx.append(int(parts[0]))
108             px.append(f16_to_float(f16_from_hex(parts[1])))
109             py.append(f16_to_float(f16_from_hex(parts[2])))
110             vx.append(f16_to_float(f16_from_hex(parts[3])))
111             vy.append(f16_to_float(f16_from_hex(parts[4])))
112             mm.append(f16_to_float(f16_from_hex(parts[5])))
113     return (np.array(idx, dtype=np.int32),
114           np.array(px, dtype=np.float64),
115           np.array(py, dtype=np.float64),
116           np.array(vx, dtype=np.float64),
117           np.array(vy, dtype=np.float64),
118           np.array(mm, dtype=np.float64))
119
120 def write_frame_host(path, idx, px, py, vx, vy, m):
121     with open(path, "w") as f:
122         f.write("# i px py vx vy m (float64)\n")
123         for i in range(len(idx)):
124             f.write(f"{idx[i]} {px[i]: .12e} {py[i]: .12e} "
125                   f"{vx[i]: .12e} {vy[i]: .12e} {m[i]: .12e}\n")
126
127 def write_frame_chip(path, idx, px, py, vx, vy, m):
128     with open(path, "w") as f:
129         f.write("# i px py vx vy m (S1E8M7 hex)\n")
130         for i in range(len(idx)):
131             f.write(f"{idx[i]} "
132                   f"{f16_to_hex(f16_from_float(px[i]))} "
133                   f"{f16_to_hex(f16_from_float(py[i]))} "
134                   f"{f16_to_hex(f16_from_float(vx[i]))} "
135                   f"{f16_to_hex(f16_from_float(vy[i]))} "
136                   f"{f16_to_hex(f16_from_float(m[i]))}\n")
137
138 def write_accel_chip(path, idx, ax, ay):
139     # accel is "chip convention" (NO GO), but written as S1E8M7 hex for chip-output compare
140     with open(path, "w") as f:
141         f.write("# i ax ay (S1E8M7 hex) [chip accel: NO GO]\n")
142         for i in range(len(idx)):
143             f.write(f"{idx[i]} "
144                   f"{f16_to_hex(f16_from_float(ax[i]))} "
145                   f"{f16_to_hex(f16_from_float(ay[i]))}\n")
146
147 # =====
148 # FP64 N-body compute (chip accel: NO GO)
149 # =====
150 def compute_acc_chip(px, py, m):
151     """
152     Compute chip-style acceleration (NO GO):
153     a_chip = sum_j m[j] * d / (r^2 + eps^2)^(3/2)
154     """
155     N = len(px)
156     ax = np.zeros(N, dtype=np.float64)
157     ay = np.zeros(N, dtype=np.float64)
158
159     for i in range(N):
160         axi = 0.0

```

```

161     ayi = 0.0
162     pxi = px[i]
163     pyi = py[i]
164     for j in range(N):
165         if i == j:
166             continue
167         dx = px[j] - pxi
168         dy = py[j] - pyi
169         r2 = dx*dx + dy*dy + EPS_SQ
170         inv3 = 1.0 / (r2 ** 1.5)
171         factor = m[j] * inv3 # <-- NO GO here
172         axi += factor * dx
173         ayi += factor * dy
174     ax[i] = axi
175     ay[i] = ayi
176     return ax, ay
177
178 # =====
179 # Main
180 # =====
181 def main():
182     os.makedirs(DIR_FRAME_HOST, exist_ok=True)
183     os.makedirs(DIR_FRAME_CHIP, exist_ok=True)
184     os.makedirs(DIR_ACCEL_CHIP, exist_ok=True)
185
186     idx, px, py, vx, vy, m = load_frame0_hex(INPUT_FILE)
187
188     # frame0 outputs
189     write_frame_host(os.path.join(DIR_FRAME_HOST, "frame0.txt"),
190                     idx, px, py, vx, vy, m)
191     shutil.copyfile(INPUT_FILE,
192                    os.path.join(DIR_FRAME_CHIP, "frame0.txt"))
193
194     for k in range(N_FRAMES):
195         # 1) chip accel (no GO)
196         ax_chip, ay_chip = compute_acc_chip(px, py, m)
197
198         # 2) write chip accel as hex (for compare with RTL chip output)
199         write_accel_chip(os.path.join(DIR_ACCEL_CHIP, f"accel{k}.txt"),
200                         idx, ax_chip, ay_chip)
201
202         # 3) host update uses K_HOST = GO * DT
203         vx = vx + ax_chip * K_HOST
204         vy = vy + ay_chip * K_HOST
205         px = px + vx * DT
206         py = py + vy * DT
207
208         # 4) write next frame (float64 host) + chip-input hex view
209         write_frame_host(os.path.join(DIR_FRAME_HOST, f"frame{k+1}.txt"),
210                         idx, px, py, vx, vy, m)
211
212         write_frame_chip(os.path.join(DIR_FRAME_CHIP, f"frame{k+1}.txt"),
213                         idx, px, py, vx, vy, m)
214
215     print("[OK] FP64 reference generated (chip-accel convention, K_HOST used).")
216     print(f"      EPS_SQ={EPS_SQ} DT={DT} GO={GO} K_HOST={K_HOST}")
217
218     if __name__ == "__main__":
219         main()

```

### A.5.5 input\_gen\_s1e8m18.py

Listing 45: code/Hardware/tb/frame\_input/input\_gen\_s1e8m18.py

```

1 # gen_input_s1e8m18.py
2 # Output 1: HEX file (S1E8M18)
3 # Output 2: DEC file (quantized decimal)
4
5 import numpy as np

```

```

6
7 # =====
8 # User knobs
9 # =====
10 N_PARTICLE = 1024
11 POS_MIN = -200.0
12 POS_MAX = 200.0
13 VEL_INIT_ZERO = True
14 VEL_MIN = -0.1
15 VEL_MAX = 0.1
16 MASS_MIN = 0.2
17 MASS_MAX = 1.0
18 RANDOM_SEED = 0
19
20 OUTPUT_HEX = "frame0_1024binit200_27bits.txt"
21
22 # =====
23 # S1E8M18 helpers
24 # =====
25 BIAS = 127
26 E_BITS = 8
27 M_BITS = 18
28 TOTAL_BITS = 27
29 E_MASK = (1 << E_BITS) - 1
30 M_MASK = (1 << M_BITS) - 1
31 SIGN_SHIFT = TOTAL_BITS - 1
32 EXP_SHIFT = M_BITS
33 U27_MASK = (1 << TOTAL_BITS) - 1
34
35 def f27_pack(sign: int, exp: int, mant: int) -> int:
36     return ((sign & 1) << SIGN_SHIFT) | ((exp & E_MASK) << EXP_SHIFT) | (mant & M_MASK)
37
38 def f27_to_hex(u27: int) -> str:
39     return f"{u27 & U27_MASK:07X}"
40
41 def f27_to_float(u27: int) -> float:
42     if u27 == 0:
43         return 0.0
44     sign = (u27 >> SIGN_SHIFT) & 1
45     exp = (u27 >> EXP_SHIFT) & E_MASK
46     mant = u27 & M_MASK
47     if exp == 0:
48         return 0.0
49     val = (1.0 + mant / (1 << M_BITS)) * (2.0 ** (exp - BIAS))
50     return -val if sign else val
51
52 def f27_from_float(x: float) -> int:
53     if x == 0.0:
54         return 0
55     sign = 1 if x < 0 else 0
56     ax = abs(x)
57
58     m, e = np.frexp(ax)
59     m2 = m * 2.0
60     e2 = e - 1
61     exp = e2 + BIAS
62
63     if exp <= 0:
64         return 0
65     if exp >= 255:
66         return f27_pack(sign, 254, M_MASK)
67
68     frac = m2 - 1.0
69     mant_f = frac * (1 << M_BITS)
70
71     mant = int(np.floor(mant_f + 0.5))
72     if mant == (1 << M_BITS):
73         mant = 0
74         exp += 1
75     if exp >= 255:
76         exp = 254
77         mant = M_MASK
78

```

```

79     return f27_pack(sign, exp, mant)
80
81     # =====
82     # Main
83     # =====
84     def main():
85         np.random.seed(RANDOM_SEED)
86
87         pos = np.random.uniform(POS_MIN, POS_MAX, size=(N_PARTICLE, 2))
88         if VEL_INIT_ZERO:
89             vel = np.zeros((N_PARTICLE, 2), dtype=np.float64)
90         else:
91             vel = np.random.uniform(VEL_MIN, VEL_MAX, size=(N_PARTICLE, 2))
92
93         mass = np.random.uniform(MASS_MIN, MASS_MAX, size=(N_PARTICLE,))
94
95         with open(OUTPUT_HEX, "w") as f_hex:
96
97             f_hex.write("# i   px   py   vx   vy   m   (S1E8M18 hex)\n")
98
99             for i in range(N_PARTICLE):
100
101                 # original float
102                 px_f = float(pos[i, 0])
103                 py_f = float(pos[i, 1])
104                 vx_f = float(vel[i, 0])
105                 vy_f = float(vel[i, 1])
106                 m_f  = float(mass[i])
107
108                 # quantize
109                 px_u = f27_from_float(px_f)
110                 py_u = f27_from_float(py_f)
111                 vx_u = f27_from_float(vx_f)
112                 vy_u = f27_from_float(vy_f)
113                 m_u  = f27_from_float(m_f)
114
115                 # write HEX
116                 f_hex.write(
117                     f"{i:4d}   "
118                     f"{f27_to_hex(px_u)} {f27_to_hex(py_u)} "
119                     f"{f27_to_hex(vx_u)} {f27_to_hex(vy_u)} "
120                     f"{f27_to_hex(m_u)}\n")
121             )
122
123         print(f"[OK] Generated:")
124         print(f"      {OUTPUT_HEX}")
125         print(f"      N={N_PARTICLE}")
126
127     if __name__ == "__main__":
128         main()
129

```