A Compilation Model for Divergent Parallel Workloads

ALEXANDER J ROOT, Stanford University, USA
CHRISTOPHE GYURGYIK, Stanford University, USA
DEVANSHU LADSARIA, Stanford University, USA
ROHAN SAWHNEY, NVIDIA, USA
JONATHAN RAGAN-KELLEY, MIT, USA
ANDREW ADAMS, Adobe, USA
MATT PHARR, NVIDIA, USA
KAYVON FATAHALIAN, Stanford University, USA
FREDRIK KJOLSTAD, Stanford University, USA

arXiv:2611.22183v1 [cs.PL] 17 Nov 2026

Efficient execution on modern parallel machines requires more than abundant parallelism; it demands instruction coherence, memory coherence, and load balance to maximize SIMD utilization and memory bandwidth. Efficiently mapping *divergent* parallel workloads to modern machines, where control flow, memory accesses, and unbounded iteration depend on per-thread runtime values, often requires global program restructuring that trades off increased coherence with reduced locality, parallelism, or extra work. We observe that the key optimizations used to accelerate divergent parallel programs can be expressed as compositions of well-known program transformations from the functional programming and loop optimization communities. We operationalize this insight through *parallel dispatch graphs*: a compilation model that represents divergent programs by representing control flow, data flow between logical stages, and parallel loops within a single abstraction. To facilitate rapid exploration of points in the optimization space, we provide a scheduling language that includes composable transformations of data-dependent control flow and data flow, together with standard loop transformations and hardware bindings. This separation of concerns allows the core algorithm to remain portable, while the execution schedule can be specialized to the specifics of target hardware and input data values. We evaluate our system using popular applications from the domain of spatial computing, and show that the optimization strategies of hand-optimized parallel tree traversals, path tracers, and Monte Carlo PDE solvers can be reproduced as schedules in our system. Our generated code matches or outperforms the hand-optimized implementations (geomeans 1.02–2.95× faster). Additionally, we demonstrate empirically that the best schedule varies across programs, machines, and inputs. Through manual exploration we find per-scene path tracing schedules that beat the expert-written schedules by a further 1.3× on the CPU and 1.6× on the GPU (geomean), up to 13×.

1 Introduction

The Single Program, Multiple Data (SPMD) programming model [17] is a (useful) lie. Programmers write scalar code for one logical thread and rely on the compiler [54] or the hardware [39] to execute the parallel program efficiently. But modern CPUs and GPUs are hierarchical *SIMD* machines that reach peak throughput only when all lanes of a SIMD unit, each running one logical thread, execute the same instruction (*instruction coherence*), access contiguous memory (*memory coherence*), and perform equal amounts of work (*load balance*); otherwise lanes sit idle and bandwidth is wasted. SPMD (purposefully) hides these facts, and for regular parallelism the lie is harmless. For *divergent* programs, whose control flow, memory accesses, and unbounded iteration depend on per-thread runtime values, it is not: mapping divergent programs efficiently onto SIMD hardware is remarkably difficult [22, 29, 35, 43], and the independent-threads abstraction falls apart.

Authors' Contact Information: Alexander J Root, Stanford University, Stanford, CA, USA, ajroot@cs.stanford.edu; Christophe Gyurgyik, Stanford University, Stanford, CA, USA, cpg@cs.stanford.edu; Devanshu Ladsaria, Stanford University, Stanford, CA, USA, devanshu@cs.stanford.edu; Rohan Sawhney, NVIDIA, Santa Clara, CA, USA, rsawhney@nvidia.com; Jonathan Ragan-Kelley, MIT, Cambridge, MA, USA, jrk@mit.edu; Andrew Adams, Adobe, San Francisco, CA, USA, andrew.b.adams@gmail.com; Matt Pharr, NVIDIA, Santa Clara, CA, USA, matt@pharr.org; Kayvon Fatahalian, Stanford University, Stanford, CA, USA, kayvonf@cs.stanford.edu; Fredrik Kjolstad, Stanford University, Stanford, CA, USA, kjolstad@cs.stanford.edu.

Many physically-based simulation methods exhibit exactly this behavior: they are massively parallel but have ample divergence [42, 52, 62, 63]. Monte Carlo path tracing [52], for example, renders an image by tracing many rays through a scene and estimating the radiance reaching each pixel. Rays that hit different materials run different shading code (instruction divergence); rays that scatter load geometry from different parts of the scene (memory divergence); and rays that bounce more times than their SIMD neighbors leave lanes idle (load imbalance). The spatial queries finding those hits, usually accelerated by tree data structures [59, 77], diverge in the same ways, as do other Monte Carlo methods [42, 62, 63], whose independent random walks diverge in trajectory and in the boundary functions they evaluate. High performance on these programs requires structuring parallelism so adjacent SIMD lanes perform the same operations on adjacent data.

Known techniques for reducing divergence gain coherence at the cost of locality, parallelism, or extra work. Wavefront path tracing [35] turns the dispatch on material type into data flow through per-material work queues, making each shading stage instruction-coherent at the cost of moving live state out of registers and through global memory. Ray reordering [40] recovers spatial (and therefore memory) coherence at the cost of sorting and synchronization. Packet tracing [76] buys memory and instruction coherence in tree traversals with extra compute. Persistent threads [3, 24] reduce load imbalance (given enough parallelism), at the cost of slightly more compute and global memory traffic. Because the *data* determines how much of each kind of divergence arises [3, 28], no single execution schedule is best across machines and inputs.

Hand-optimized path tracers [16, 37, 52], Monte Carlo PDE solvers [62, 63], and traversal libraries [77] each compose several of these techniques with standard loop optimizations such as tiling and vectorization [8, 54] and with explicit binding to hardware accelerators [39, 47]. However, these systems can express that composition only by writing it into the application: the programmer no longer writes one logical thread but the SIMD-level structure of a specific machine. Each system is thus tied to the hardware it was tuned for (NVIDIA GPUs [52], AMD GPUs [41], x86 vector CPUs [37]), and exploring a different schedule means restructuring the application.

Recent programming systems make orthogonal steps in productive directions. Luisa [81] and Dr.Jit [28] maintain an SPMD programming model but let the user instruct the compiler to split a renderer into stages for better coherence. Luisa splits the program at user-annotated suspension points and lets the user pick from a fixed set of schedulers (e.g., wavefront [35] or persistent threads [24]), while Dr.Jit automatically decides suspension points for wavefront execution. These are coarse-grained, global transformations, however, and neither compiler exposes other scheduling commands to compose them with. User-schedulable languages (USLs) similarly provide data-parallel frontends, along with a scheduling language of composable loop transformations [15, 26, 55, 65] and hardware bindings [9, 26]. However, USLs have primarily been developed for feed-forward array programs and lack scheduling commands targeted at divergence-mitigation.

Our key insight is that the hand-optimized implementations of divergent parallel programs [37, 52, 60–63, 77] can be automatically generated via a directed composition of well-known program transformations from the functional programming [7, 58] and loop optimization [8, 54] communities. We propose *parallel dispatch graphs* (PDGs), a compilation model that represents control flow, data flow, and parallel loops in a single abstraction, enabling these transformations to compose. PDGs mirror how graphics researchers describe their systems, as pipelines of stages connected by work queues [35, 52, 72], but are precise enough to serve as an intermediate representation. A PDG models a program as a hierarchy of parallel loops, each draining either the iteration space of a programmer-written parallel for-loop or a work queue induced by scheduling transformations that turn control flow into data flow. Loop bodies are control-flow graphs whose edges are either *eager*, transferring control immediately, or *deferred*, pushing a continuation onto a queue to run when that queue is drained. Our model preserves the useful lie of SPMD for programmers by pairing a

machine-agnostic SPMD frontend with a scheduling language that specializes the program to its hardware and data. We make the following contributions:

- Parallel dispatch graphs (PDGs), an abstraction that models divergent parallel execution as a pipeline of logical stages within a hierarchy of parallel loops.
- A push-based scheduling language of composable PDG transformations: divergence mitigation [35, 81], standard loop transformations [8], and binding to hardware resources [9, 26].

We implement these ideas in Grove, a parallel extension of the Bonsai [25, 59] compiler for portable tree-accelerated spatial queries. Grove uses PDGs as an intermediate representation and compiles to x86 (SSE4/AVX2/AVX-512) and NVIDIA GPU backends. Our evaluation shows that Grove reproduces the state-of-the-art optimizations of parallel tree traversals [60, 77], path tracers [37, 52], and Monte Carlo PDE solvers [62, 63] as distinct points in one scheduling space, matching or outperforming the hand-optimized originals (geomean speedups of 1.02–2.95×). With Grove we find path tracing schedules that beat the expert schedules by a further 1.3× on the CPU and 1.6× on the GPU (geomean), up to 13×, and show the best schedule is often *data-dependent*, motivating a scheduling language for rapid exploration of the execution space.

Scope and Limitations. Grove targets divergent workloads such as parallel tree traversals and parallel random walks, as opposed to *irregular* workloads, where divergence comes from irregularity in the data structures, such as sparse tensor algebra [31, 71], relational algebra [19], and graph computations [79]. Grove does not support irregular parallel loop nests. Additionally, Grove does not verify user-provided parallelism annotations; a lack of independence across parallel iterations is undefined behavior. Grove does not yet have an autoscheduler.

2 Optimizing a Simple Path Tracer

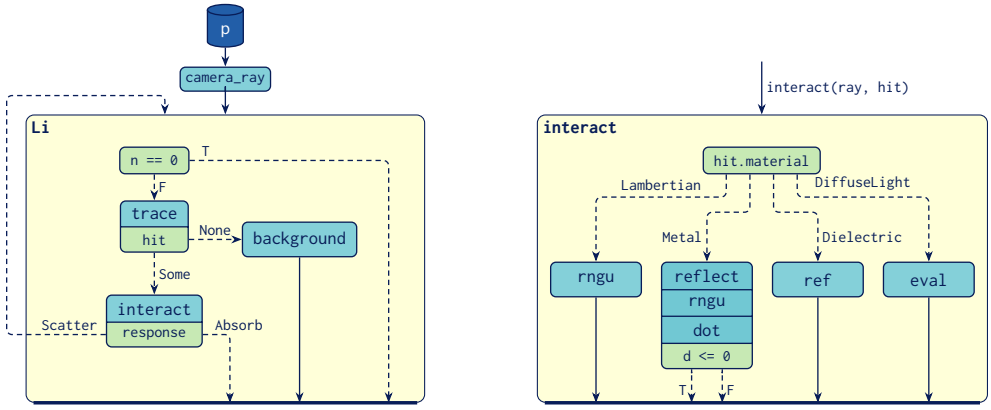
Consider the simple path tracer in Listing 1, which fills an output image `out` by sending a ray per pixel into a geometric scene and computing the illumination returned along the ray to the pixel. A path tracer has ample parallelism, because every pixel performs independent work, and is cleanly written in SPMD style: `render` is a parallel loop over pixels, and the loop body is scalar code for one pixel. However, a path tracer also has ample divergence. `Li` first calls `trace`, which loads nearby data from the scene data structure for rays traveling in similar directions [76] but loses that spatial

```
func interact(ray, hit) → Response:
  match hit.material:
  | Lambertian(albedo) →
    return Scatter{Ray{hit.pt, hit.n + rngu()}, albedo, 0}
  | Metal(albedo, fuzz) →
    dir = reflect(ray.dir, hit.n) + fuzz * rngu()
    if dot(dir, hit.n) <= 0: return Absorb{0}  // grazing
    return Scatter{Ray{hit.pt, dir}, albedo, 0}
  | Dielectric(eta) → // reflect or refract
    return Scatter{Ray{hit.pt, ref(ray.dir, hit, eta)}, 1, 0}
  | DiffuseLight(tex) → return Absorb{eval(tex, hit.pt)}

func Li(L, beta, ray, n, scene) → Value:
  if n == 0: return L
  else match trace(ray, scene): // Find first hit geometry in scene
  | None → return L + beta * background(ray)
  | Some(hit) → match interact(ray, hit):
    | Absorb(e) → return L + beta * e
    | Scatter(dir, a, e) → return Li(L + beta * e, beta * a, dir, n-1, scene)

func render(width, height, out : mut, depth, scene):
  parfor p in [0, width * height): // compute pixel p
    out[p] = Li(0.0, 1.0, camera_ray(p), depth, scene)
```

Listing 1. A *Ray Tracing in One Weekend* (RTIOW) path tracer [66] with a single sample per pixel, written in Grove. ADT types for materials and interaction responses, and argument types, are omitted for brevity.



(a) `render` and `Li` of Listing 1; calls are not expanded. (b) `interact`: dynamic dispatch on material type.

Fig. 1. Parallel dispatch graphs (PDGs) representing divergence in the RTIOW path tracer. Blue queues are parallel loops; teal boxes are (unexpanded) calls; green boxes with dashed arrows are dispatches (conditional control flow); solid arrows are unconditional control flow; arrows to the bottom are returns.

coherence after the first scattering step [2, 53] (memory divergence). The **`match`** on the result of `trace` calls `background` for rays that miss geometry and `interact` for rays that hit geometry, and `interact` applies a different shading method per material (instruction divergence). The **`match`** on the result of `interact` recursively evaluates illumination for scattering rays, while absorbed rays return an illumination value (load imbalance).

Figure 1 illustrates the parallel dispatch graphs (PDGs) of `Li` and `interact`: a parallel loop over pixels executes the control flow of `Li`, which calls functions (teal boxes) and dispatches on sum types (green boxes). The dashed arrows are the possible sources of divergence.

Consider a SIMD trace of the functions called within `Li` across four lanes, read from left to right:

Lane 0	trace	background					
Lane 1	trace		interact	trace		interact	trace	background
Lane 2	trace		interact	trace	background			
Lane 3	trace		interact	trace		interact	trace	background

Each white space, or bubble, is instruction (or occupancy) divergence, where a lane is masked out as another executes, and represents wasted work. Optimizing this program means minimizing the bubbles, weighted by cost (e.g., thousands of cycles for a bubble alongside `trace`).

A useful transformation *defers* execution of `background` until all lanes have finished tracing and shading, which results in the execution trace on the right. For just these four lanes of work, this trace is optimal:

Lane 0	trace					background
Lane 1	trace	interact	trace	interact	trace	background
Lane 2	trace	interact	trace			background
Lane 3	trace	interact	trace	interact	trace	background

the long-pole lanes have no bubbles. However, the dispatch graph must be fundamentally rewired to defer execution of `background`. Figure 2 shows this rewiring: the control flow that originally executed `background` becomes *data flow* into the queue `b`, which saves the (continuation) state needed to resume `background`, and a second parallel loop *drains* `b` and executes `background`. Listing 2 shows the code Grove generates for the transformed PDG. Note that the change is nonlocal:

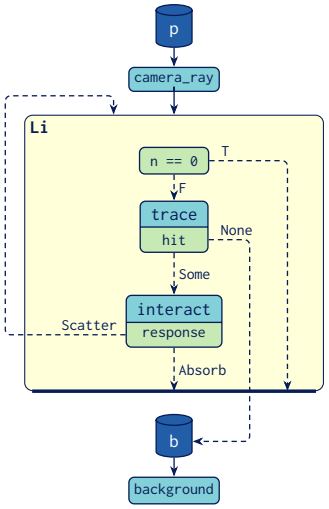


Fig. 2. PDG with background deferred.

```
// Li must communicate whether work was deferred
type Result = | Done(v : Value) | Defer(i : Index);

func Li(L, beta, ray, n, scene, b : mut) → Result:
 if n == 0: return Done(L)
 else match trace(ray, scene):
 | None → return Defer(append(b, {L, beta, ray}))
 | Some(hit) → match interact(ray, hit):
  | Absorb(e) → return Done(L + beta * e)
  | Scatter(dir, a, e) →
    return Li(L+beta*e, beta*a, dir, n-1, scene, b)

func render(width, height, out : mut, depth, scene):
 b = queue()
 parfor p in [0, width * height):
  ray = camera_ray(p)
  match Li(0.0, 1.0, ray, depth, scene, b):
  | Done(v) → out[p] = v
  | Defer(i) → b[i].p = p // finish continuation

 // Drain the queue
 parfor (L, beta, ray, p) in b:
  out[p] = L + beta * background(ray)
```

Listing 2. Grove-generated code for Figure 2: Listing 1 with background deferred into a global work queue.

the value of p was not accessible within Li, requiring render to patch p into the queue element to complete the continuation. The code changes considerably, but the PDG differs in one edge and the induced drain: every edge in Figure 1 is *eager* and transfers control immediately, while the edge into background in Figure 2 is *deferred*, saving a continuation into b for the drain to execute later.

While deferral is useful in its own right, production renderers compose deferral with other transformations. The PBRT GPU renderer [52] tiles the loop over pixels so that a single set of queues fits in device memory, defers many stages (e.g., material shaders and shadow ray queries) into work queues in the wavefront style of Laine et al. [35], specializes the shading stage per material type, and accelerates ray tracing and texture evaluation through hardware primitives [45, 47]. Figure 3 visualizes PBRT's execution model applied to the simple renderer: each transformation is a local edit to Figure 1, and the recursion of Li has become a cycle through the trace queue t. Listing 3 shows the generated code: six parallel loops over queues and a loop that drains t until empty, which bears little resemblance to the clean SPMD specification of the algorithm. The next section presents the conceptual model behind these transformations, showing how PDGs capture the structure of divergent parallel execution, and how scheduling commands transform a PDG.

3 Scheduling Model: Parallel Dispatch Graphs

Parallel dispatch graphs model the *structure* of divergent parallel execution: a hierarchy of parallel loops whose bodies are control-flow graphs with eager and deferred edges, where each loop drains the iteration space of a programmer-written **parfor** or a schedule-induced queue. Scheduling maps this graph onto the structure of the machine. We first describe the components of a parallel dispatch graph, then the three classes of rewrite a schedule composes. The next section describes the programming system built on this model.

3.1 Components of a Parallel Dispatch Graph

A parallel dispatch graph has four transformable components:

Parallel loops drain either the range of a programmer-written **parfor** (Figure 1 loop p) or a schedule-induced queue (Figure 2 loop b). Nested loops form hierarchy.

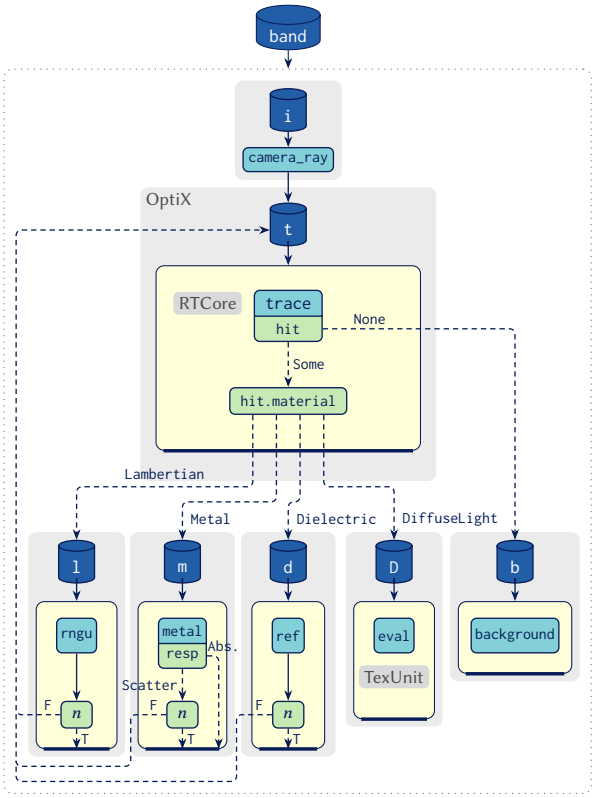


Fig. 3. Wavefront PDG: the continuation of `Li` splits hits into per-material queues; scattering materials push into the trace queue; hardware acceleration is labeled in gray, unlabeled gray backgrounds map to standard GPU kernels.

```
func render(width, height, out, depth, scene):
 for band in [0, (width*height+(TILE-1))/TILE):
  t, b, l, m, d, D = queue() × 6
  parfor i in [0, TILE):
   p = band * TILE + i
   if p < width * height: // assumes depth > 0
    append(t, {0, 1, camera_ray(p), depth, p})
  while !empty(t): // loop induced by back arrow
   parfor (L, beta, ray, n, p) in t: // OptiX
    match trace(ray, scene): // RTCore
    | None → append(b, {L, beta, ray, p})
    | Some(hit) → s = {L, beta, ray, hit, n, p}
       match hit.material: // material binning
       | Lambertian(_) → append(l, s)
       | Metal(_) → append(m, s)
       | Dielectric(_) → append(d, s)
       | DiffuseLight(_) → append(D, s)
   // interact, specialized per material (queue)
   parfor (L, beta, ray, hit, n, p) in l:
    dir, a = lambertian(ray, hit) // Scatter
    if n - 1 > 0:
      append(t, {L, beta*a, dir, n-1, p})
    else: out[p] = L
   parfor (L, beta, ray, hit, n, p) in m:
    match metal(ray, hit):
    | Absorb(e) → out[p] = L + beta * e
    | Scatter(dir, a, _) →
      if n - 1 > 0:
        append(t, {L, beta*a, dir, n-1, p})
      else: out[p] = L
   parfor (L, beta, ray, hit, n, p) in d:
    dir, a = dielectric(ray, hit) // Scatter
    if n - 1 > 0:
      append(t, {L, beta*a, dir, n-1, p})
    else: out[p] = L
   parfor (L, beta, ray, hit, n, p) in D:
    e = light(ray, hit) // on tex. unit, Absorb
    out[p] = L + beta * e
   parfor (L, beta, ray, p) in b:
    out[p] = L + beta * background(ray)
```

Listing 3. Grove-generated code for Figure 3: a wavefront version of Listing 1 in the style of PBRT [52]. `parfor`s are on GPU threads, and iteration over a queue resets the size to 0.

Calls transfer control to a function, whose body is a control-flow graph of its own (the teal boxes of Figure 1). A call, or its continuation, can be deferred.

Dispatches on a sum type, including booleans, are the only conditional control flow (the `match` on `hit.material` in Figure 1b). A dispatch is *varying* if iterations of the enclosing loop can take different arms and *uniform* if all take the same arm; only varying dispatches produce SIMD divergence. Dispatch edges can also be deferred.

Edges are *eager* or *deferred*. An eager edge transfers control immediately. A deferred edge is *data flow*: the edge appends the live state of the iteration and a continuation to a queue, whose drain executes the continuation (the edge into b in Figure 2).

Unconditional control flow (jumps and returns) is coherent by nature and not manipulated by the schedule. Before scheduling, a parallel dispatch graph is a control flow graph with parallel annotations, as in other parallel intermediate representations [57, 64]; ours differs in exposing data flow as a schedulable construct.

Every parallel loop, and therefore every queue, is *owned* at a level of the hierarchy, a loop or the root of the computation, and the drain of a queue is placed at the owner, after every producer. For

example, the root owns b in Figure 2, and the band loop owns every queue in Figure 3. A queue may have several producers with different continuations, in which case the stored states carry a continuation tag that the drain dispatches on. A drain can also feed a queue upstream of itself, which happens when a schedule defers part of a recursive function; the resulting cycle executes by iterating until the queue is empty, as the `while !empty(t)` loop of Listing 3 does for the backedge into `t` in Figure 3. Queues that directly feed themselves are double-buffered (Section 8).

3.2 Scheduling as PDG Rewriting

Every scheduling command is a *transformation* of the PDG. Grove supports three classes of transformations: loop transformations that rewrite the hierarchy, control flow transformations that mitigate divergence and specialize code, and bindings that attach execution to hardware resources.

Loop Transformations. Loop transformations create or remove hierarchy to match the hierarchy of the machine: tiling splits a loop in two, collapsing and interchange flatten and reorder nested loops, and a *persistent* transformation replaces a parallel loop with a loop of constant extent whose iterations claim the original iterations from an atomic counter to load balance [3, 24]. These transformations apply to any parallel loop, whether the loop drains a `parfor` or a queue.

Control Flow Transformations. Control flow transformations mitigate divergence and specialize code. *Deferral* replaces an eager edge with a deferred edge into a queue, capturing the continuation [7, 81]; the drain then executes the continuations together, at the cost of writing the continuations to memory and reading them back. Deferral is nonlocal: the continuation captures the live state of every frame between the owner and the call, and each frame on that path is rewritten to report whether work was deferred, as `render` patches `p` in Listing 2. *Specialization* splits a PDG (or sub-PDG) into one copy per variant of a sum type, so that each copy runs a single arm of the dispatch on that type. On a uniform dispatch, specialization is loop unswitching [4]: the dispatch moves out of the loop, which reduces register pressure and lets each arm take a different schedule. On a varying dispatch, specialization moves the dispatch into the data flow: each iteration is routed to a per-arm queue, and each drain runs one arm alone, with no divergence at that dispatch. *Sorting* orders the iterations of a loop by a key, so that neighboring lanes take the same arms or touch nearby memory [40, 46]. *Loopification* turns a (linear or tree)-recursive function into a loop with an explicit stack [68].

Binding to Hardware. The final step of scheduling binds the PDG to the machine. Parallel loops can be bound to a parallel resource (e.g., CPU threads or GPU blocks or threads), functions can be bound to fixed-function accelerators (RT cores [45], texture units [47]), and queues can be assigned to a level of the machine's memory hierarchy. Bindings must respect the hierarchy of the machine (Section 6.3), e.g., a loop bound to GPU threads must lie beneath one bound to GPU blocks. For example, in Figure 3, `trace` can be assigned to the RT cores because the drain of `t` is bound to OptiX threads [48], a requirement of the NVIDIA programming model for RT cores. Bindings are considered *final*, meaning Grove rejects any further transformation of a bound loop or call.

Vectorization. Vectorization is a loop transformation, a control flow transformation, and a binding at once: it binds the iterations of a parallel loop to hardware vector lanes and linearizes the varying control flow within the loop [43, 54], preserving uniform control flow.

Scheduling Example. Listing 4 composes all three classes to transform Figure 1 into Figure 3, and therefore Listing 1 into Listing 3, in seven steps:

(1) Tile the `p` loop into an outer loop `band` and an inner loop `i`.
(2) Defer evaluation of `trace` into a work queue, `t`, owned by an iteration of `band`.
(3) Defer evaluation of `interact` into a queue also at `band`.
(4) Specialize the `interact` work queue on material type, producing work queues (`l`, `m`, `d`, `D`).

(5) Defer computation of `background` into a sixth work queue, `b`, at band.
(6) Bind the drain of `t` to OptiX [48] threads, the evaluation of `trace` to RT Cores [45], and the evaluation of `eval` to the texture units [47].
(7) Tile and bind the `i` loop and drains of the remaining queues to GPU blocks and threads.

The remaining band loop is not bound to a parallel resource, so it runs sequentially.

```
band, i = render.split(p, TILE) // split p by TILE into an outer band loop and inner i loop
t = render.queue(band, Device) // create a queue in Device memory inside of a band iteration
k0 = Li.defer(trace, t) // defer execution of trace into queue t, continuation is `k0`
q = render.queue(band, Device) // create a queue in Device memory inside of a band iteration
k1 = k0.defer(interact, q) // defer evaluation of interact into queue q
l, m, d, D = k0.specialize(q, hit.material) // q splits per sum arm of Material ADT
b = render.queue(band, Device) // create a queue in Device memory inside of a band iteration
k2 = k0.defer(background, b) // defer execution of background into queue b
render.bind(t, OptixThread) // bind the drain of t to OptixThreads
k0.bind(trace, RTCore) // bind evaluation of trace within k0 to the RT Core
interact.bind(eval, TextureUnit) // bind evaluation of eval to the texture unit
render.gpu_tile(i, 256) // syntactic sugar for x, y = split(i, 256); bind(x, GPUBlock);
                        //                           bind(y, GPUThread)
// also gpu_tile l, m, d, D, and b
```

Listing 4. The schedule that transforms Figure 1 into Figure 3, in the scheduling language of Section 6.

4 Programming System Overview

We build Grove as a parallel extension to Bonsai [25, 59], a strict, statically-typed functional query language that produces efficient tree traversals for spatial queries. Tree traversals are themselves divergent programs [3], and are used in divergent parallel spatial Monte Carlo simulators like path tracing and spatial PDE solvers [42, 62, 63]. Thus, a compiler that abstracts away the tree data structures [59], their memory layouts [25], and now their parallel execution strategies lets programmers productively explore how to optimize spatial query programs.

Figure 4 illustrates how Grove extends the Bonsai compilation pipeline. We inherit the set operation and algebraic data type (ADT) specification languages of Bonsai [59] and the Scion memory layout specification language [25]. We extend the frontend programming model with SPMD-style *parallelism annotations*, whose syntax and semantics are described in Section 5. Section 6 describes the syntax of our scheduling language, in which transformations consume and produce code *cursors* [27], pointers to code objects.

Section 7.2 describes the simple lowering of Grove's SPMD frontend to our PDG intermediate representation (Section 7.1), which runs after Bonsai has lowered set expressions to explicit tree traversals [59] and before it specializes memory layouts [25]. The schedule is then applied to the

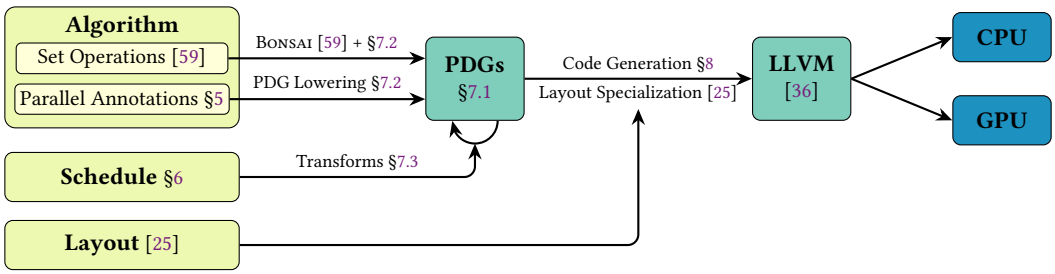


Fig. 4. Grove compilation pipeline. Set operations and parallel annotations are lowered to PDGs, which the schedule rewrites. Code generation and layout specialization perform some vector instruction selection (among other lowering optimizations) before using LLVM to target CPUs and GPUs.

PDG as a series of transformations (Section 7.3). Lastly, Section 8 describes code generation from the transformed PDG to LLVM IR [36], including layout lowering [25] and vector instruction selection.

5 Expressing Parallel Algorithmic Logic

Grove extends Bonsai with the imperative structure of SPMD programming interfaces [39, 54] (Figure 5): standard control flow (**if**/**match**/**for**) and recursion, which Bonsai lowers set operations to [59]; a parallelism annotation, **parfor**; and access-restricted types. Grove is thus a multi-paradigm language: parallel structure is exposed imperatively, while spatial queries are expressed as first-class operations over sets [59]. **parfor** asserts that a loop's iterations are independent, expressing the *possibility* of parallelism without the guarantee of parallel execution; **for** and **while** loops are sequential, and loop transformations apply only to parallel loops.

D : Decl ::= **func** $f(\overline{x : T}) \to T\{\ S^*\ \}$
| **type** $E = (\mid V(\overline{x : T}))^+$; ADTs
T : Type ::= t | **mut** t | **reduce**$(\oplus)$ t
S : Stmt ::= $x : T = e$;
| [**atomic**] $x \oplus\!= e$; reductions
| **parfor** i **in** $e_{lo} : e_{hi}[: e_s]\{\ S^*\ \}$ independent work
| **for** i **in** $e_{lo} : e_{hi}[: e_s]\{\ S^*\ \}$ sequential
| **if** $e_b\{\ S^*\ \}$ [**else**$\{\ S^*\ \}$] sugar for **match**
| **match** $e\{\ (\mid V(\overline{x}) \to S)^+\ \}$
| **while** $e_b\{\ S^*\ \}$
| **return** [e] ; | $g(\overline{e})$;
e : Expr, t : type ::= ... Bonsai [59]

Fig. 5. Grove's extensions to Bonsai [25, 59].

The programmer is responsible for the correctness of their parallel annotations: Grove does not verify that **parfor** iterations are independent. Instead, inspired by Liszt [18], we restrict program behavior to help the programmer write safe parallel code: all data is read-only, reduce-only, or write-only. Aliasing is undefined behavior. The type system enforces the access restrictions by making data immutable by default, with opt-in `mut` types for write-only data and `reduce`$(\oplus)$ types for reduce-only primitives; reductions may also be marked **atomic**.

6 Schedule Language Syntax

A Grove schedule is a sequence of directives that transform a program's PDG, defining an execution order and parallelization strategy without changing *what* is computed. Each directive applies a transformation of Section 3.2 to a *cursor* (Section 6.1), a pointer to a piece of a PDG; some directives construct new cursors (e.g., loop splitting returns two loop cursors, given one). Directives that bind to hardware are checked against a hardware model (Section 6.3).

6.1 Cursors

Cursors [27] represent pointers to specific pieces of code. In Grove, cursors are loop, function, or edge references and can come from three places: program cursors from the original program (**parfor** loop indices, functions, and variables within functions), schedule-induced cursors constructed by a scheduling command, and data-structure-induced cursors, where the lowering of a Bonsai [25, 59] tree traversal produces a loop (e.g., over the children of a tree node).[1] Listing 4 uses program and schedule-induced cursors, and Listing 6 in Section 9.2 uses data-structure-induced cursors to schedule tree traversals. Figure 6a includes syntax for our cursor notation, which includes loops, calls, call continuations (`g.cont`), and match arms (`x::V` for the `V` arm of the match on variable `x`).

6.2 Transformations

Figure 6a gives the syntax of Grove's scheduling language. Figure 6b gives each transformation's signature, indicating which produce new cursors.

Loop transformations are the standard transformations seen in Halide [55, 56] and Exo [26, 27], with one exception. `persistent(l, n)` is the persistent-thread transformation of Section 3: `n`

[1] If tree traversals were written directly, instead of using Bonsai, these would just be program cursors.

S : Schedule ::= **schedule** { D^* }	
D : Directive ::= [c (, c)* =] (f | q) (.(L | K | H))+	
c : Cursor ::= l | f | q | e	handles into a PDG
l : Loop cursor ::= i | x::V.a	
e : Edge cursor ::= g | g.cont | x::V	call, call return, arm
o : Owner ::= l | root	a loop, or loop nest root
L : Loop ::= split(l, n)	tiling loop l by a factor of n
| collapse(l_o, l_i)	loop flattening
| reorder(l_i, l_o)	loop interchange
| persistent(l, n)	n iterations claim l iterations
| unroll(l)	loop unrolling
| vectorize(l)	vector lane per iteration
K : Control ::= queue(o [, M] [, n])	queue owned at o
| defer(e, q)	continuation at e onto q
| specialize(c, x)	execution per variant of x
| sort(c, λ)	instances of c ordered by λ
| loopify([n] [, M])	recursion to a loop; stack n
H : Hardware ::= bind(c, R)	one R per instance of c
| prefetch(l [, n])	n cache lines prefetched

(a) Syntax. f, g name functions, q a queue, i a loop, x a variable, V a variant, a a field, n an integer, and R, M a resource and storage kind (Section 6.3); x::V.a is a data-structure-induced loop over V's field a.

Args, *Key*, *Var*	Instance arguments, sort key, variable
Resource, *Storage*	Machine-dependent, Section 6.3
Cursor = *Node* | *Edge*	Handle into a PDG
Node = *Loop* | *Func*	Loops and calls
Queue <: *Loop*	Queue drain is always a loop
Owner = *Loop* | root	Queue owners
split	: *Loop* × *Int* → *Loop* × *Loop*
collapse	: *Loop* × *Loop* → *Loop*
reorder	: *Loop* × *Loop* → ()
persistent	: *Loop* × *Int* → *Loop*
unroll	: *Loop* → ()
vectorize	: *Loop* → ()
queue	: *Owner* [×*Storage*] [×*Int*] → *Queue*
defer	: *Edge* × *Queue* → *Func*
specialize	: *Node* × *Var* → *Node* × ⋯ × *Node*
sort	: *Node* × (*Args* → *Key*) → ()
loopify	: [*Int*][×*Storage*] → ()
bind	: *Node* × *Resource* → ()
prefetch	: *Loop* [×*Int*] → ()

(b) Signatures. Each transform is a method of the cursor it is applied within: a function, or a queue for its drain. Result types are the cursors produced by the transform.

Fig. 6. Grove's scheduling language. Directives apply in written order. [...] marks an optional argument/type.

iterations claim the iterations of `l` from an atomic counter. Composed with `split`, the iterations claim chunks instead, which reduces contention on the counter, as proposed by Aila and Laine [3]. Additionally, `split` has a three-loop form (omitted for brevity) that returns a third cursor for the *tail loop* when the split is not exact.

Control flow transformations are the novel scheduling language contribution of Grove. Transforming control flow into data flow is composed of two separate transformations:

Queueing `q = f.queue(i[, M][, n])` places a queue `q`, owned at the loop `i` of `f`, in storage `M` with capacity `n`; `f.queue(root)` places the queue outside of any loop. `M` must be available at `i` in the hardware model (Section 6.3). If `n` is omitted, Grove infers it (Section 8).

Deferral `c = f.defer(e, q)` rewrites the edge `e` within `f` into an append to `q`, which must be owned within `f` or within a caller of `f`. The returned cursor `c` is the continuation the drain of `q` executes.

If several `defer` directives target one queue, the drain dispatches to their continuations (Section 7.3).

Other control flow transformations include: `specialize(c, x)`, which splits the cursor `c` per variant of the ADT `x` and returns one cursor per variant (e.g., the per-material queues in Listing 4), and can produce unswitched loops for uniform control flow (Section 7.3). `sort(c, k)` orders the instances of `c` by the key `k`: e.g., the rays in a queue by a spatial hint before tracing or the recursive calls of a tree traversal (Section 9.2). Lastly, `loopify([n][, M])` turns a tail-recursive function into a loop [68]; for tree recursion, `n` and `M` size and place a software stack (Section 9.2).

Finally, Grove supports hardware binding in the style of Exo [26] and Mosaic [9]. `bind(c, R)` executes the cursor `c` on the resource `R`: a loop on CPU threads or GPU blocks or threads, or a function on an accelerator such as the RT cores [45] or texture units [47]. `prefetch(l, n)`

prefetches `n` cache lines ahead of the loop `l`, as in Halide [55]. `gpu_tile(l, n)` is syntactic sugar for `x, y = split(l, n); bind(x, GPUBlock); bind(y, GPUThread)`.

6.3 Hardware Models

A *hardware model* describes a machine as a hierarchy of compute resources R and the storage kinds M each level owns. `bind`, `queue`, and `loopify` are checked against the model in lowering: storage of kind M is owned at a level that lists M and used only at or beneath its owner. For example, an `RTCore` binding must lie within an `OptixThread` loop [45, 48], and a `DeviceShared` queue within the block that owns it. `Device` is the exception: the host owns it, as band owns the queues of Listing 4, but only GPU loops use it. The design follows Sequoia [21] and Exo [26], and can easily extend to other machines.



Fig. 7. Grove's two machine models: compute resources R in their hierarchy and storage kinds M at the level that owns them.

7 Parallel Dispatch Graph Intermediate Representation

We present PDGs as an SSA intermediate representation (IR), outline the lowering from Grove's frontend to PDGs, and show how schedule transformations are expressed as IR rewrites.

7.1 SSA Blocks with Terminators

We implement PDGs as an SSA intermediate representation with explicit block arguments [6]. Block terminators represent all control flow, parallelism, and data flow, and explicitly pass continuation parameters. The live state at a call is therefore the argument list of the continuation block, and continuation capture needs no liveness analysis [30]. Most transformations (other than vectorization) solely manipulate the terminators of blocks, as Section 7.3 will show. Instructions within a block are arithmetic, side-effecting operations such as stores, and explicit vector operations, as in LLVM [36] and Halide [55, 56]. Grove's IR has 8 terminators:

Jump `jmp` terminators are standard unconditional jumps. They pass arguments from the source block's scope to the target block.

Conditional Jump `cjmp` terminators are n-ary conditional jumps on a sum type (the dispatches of Section 3) and the only conditional control flow in the IR. The variables bound by each arm are prepended to the arguments of that arm's block.

Call `call` terminators transfer control flow to a function target, also expressing the return continuation block. If the function has a return type, the return value is prepended to the continuation's argument list.

Return `ret` terminators return control flow to the caller, optionally returning a value.

Multicall `mcall` terminators are a sequence of calls to a function, representing Bonsai's `from` node [59]: unordered tree recursion on a static set of arguments. This node simplifies some pattern matching in Section 7.3.

Parallel `par` terminators take an index name, an iteration slice (start/end/stride), a loop body block (which takes the loop index as the first argument), and a continuation block. The arguments to the loop body block (other than the index) are *uniform* arguments.

Yield `yield` terminators end parallel loop bodies and take no arguments.

Defer `defer` saves a continuation to a named queue and then yields control flow, returning the queue index if the defer is nested within a callee.

Listing 5. PDG block terminators; optional fields are marked ?.

```
jmp block(args)
cjmp v [block0(args0), ..., block{n-1}(args{n-1})]
call f(argsf) cont(argsb) hw_tag?
ret val?
mcall f [args0, ..., args{n-1}] mask? order? cont(argsb)
par idx range body(argsu) cont(argsc) hw_tag?
yield
defer args queue
```

Listing 5 summarizes these terminators. In listings, a continuation written as **ret** or **yield** is an empty block ending in that terminator. Note that any results of parallelism must be communicated via explicit memory stores and atomics: **par** does not return a value.

7.2 Lowering to PDGs

Lowering an input program into a PDG is straightforward, via the SSA construction of Braun et al. [11]. **match** lowers to **cjmp**, as does any conditional control flow (**if** is a **cjmp** on the condition); calls and returns lower to **call** and **ret**. **parfor** lowers to **par**, with the body of the loop as the body block and the code after the loop as the continuation. Bonsai's `from` lowers to **mcall**. Listing 7 in Appendix A shows the lowering of Listing 1 after simplification.

7.3 Transforming PDGs

Loop Transformations. Loop transformations other than vectorization rewrite the **par** terminator the cursor points to and leave the loop body alone [8, 55]. Vectorization also rewrites the body: we implement Moll and Hack's partial control-flow linearization [43], which provably preserves uniform control flow where possible.

Queue. `queue` places a symbolic allocation at the cursor. The queue records every `defer` that pushes into the queue. The `defer` transformation places the drain and Section 8 sizes the allocation.

Defer. Deferral, `c = f.defer(e, q)`, is a selective CPS conversion [7] of the call path from the loop that owns `q` down to `e`, followed by defunctionalization [58] of the continuation into the state stored in `q`. Instead of executing the edge `e`, the producer stores the live state of every frame on the path into `q` and ends execution. The drain of `q` loads that state and executes the deferred control flow. Every function on the path to the deferred edge either returns its value (`Done`) and executes normally, or the index of the stored state (`Defer`) and saves the live state within the frame at that index. In Listing 2, the path runs from the loop over `p` through `Li` to `background`: `Li` stores `L`, `beta`, and `ray` in `b` and returns the index, `render` adds `p`, and the drain of `b` finishes both functions. In the IR, `e` becomes a **defer** of its arguments and its continuation's arguments, and the body of the owning loop yields. The live state of a frame is the arguments of its continuation block, except cheap (and pure) recomputable values [12]. Tail calls leave no live state, e.g., the recursive call in `Li`. Non-tail recursive calls on the path require unbounded state; Grove rejects defer commands that cannot be statically bounded. The drain is placed at the owner, after all producers, and is a parallel loop over `q` that executes `e` and the remainder of the frames of the continuation; this drain code is the cursor `c`. A queue with several producers stores a continuation tag with each state, on which the drain dispatches; a function that defers on every path drops `Done` and its dispatch, as in Listing 3. Algorithm 1 in Appendix B gives the full algorithm for this transformation.

Grove's type restrictions (Section 5) make legality simple to check: deferral is legal as long as no statement between a producer of a queue and its drain changes the access type of data the deferred continuations touch (e.g., reads data the deferred edge writes). Within a **parfor**, iterations are independent and each datum is read-only, write-only, or reduce-only, so pieces of its iterations may run in any order (up to order of reductions); only a statement after the loop can read what an iteration of the loop wrote. Grove therefore checks every statement that can run between a producer and the drain against the data the continuation accesses: each must maintain the access type it has in the producer, otherwise the transformation is rejected.

Specialize. Specialization takes a cursor and an ADT variable in scope at that cursor, and returns one copy of the cursor per variant. `l, m, d, D = k0.specialize(q, hit.material)` in Listing 4 splits the appends into `q` by material and the drain of `q` into per-material drains, a specializing defunctionalization [10, 58]. For a loop and a uniform variable, specialization is loop unswitching [4]. In our implementation of PBRT's renderer, for example, the parallel loop over pixels dispatches on the integrator that computes radiance; specializing on the integrator ADT yields one parallel loop per integrator, each with a different schedule.

Sort. Sorting has different lowerings for different cursors. Sorting a loop over a queue lowers to vendor-provided sort implementations [14, 75]. On a function within a loop body, sort reorders the iterations of the loop body, e.g., via OptiX's shader execution reordering [46]. On an **`mcall`**, recursion into a tree's children [25], the sort fixes the order in which the children are visited, which is always legal for traversals lowered by Bonsai [59], and lowers to a sorting network [1].

Loopify. Loopify applies to a tail-recursive function, linear or tree. For linear recursion, the standard tail-call optimization [68] rewrites **`call`** `f(args)` **`ret`** into **`jmp`** `entry(args)`. For tree recursion, the recursion Bonsai produces [59], loopification places an explicit stack of the given size at the entry of the function and rewrites **`mcall`** `f [...]` **`ret`** into pushes onto the stack followed by a jump to the loop head.

Bindings. Bindings attach hardware tags to function calls and loop bodies in the PDG. Grove reports an error if the target is already tagged or if a further transformation is applied to the cursor.

8 Code Generation and Optimization

Grove applies standard optimizations [32, 33, 44, 78] before lowering to LLVM IR [36] by translating block arguments to ϕ-nodes. We describe the non-standard parts of code generation below.

Layouts. The final lowering to LLVM also stamps in concrete memory layouts for tree data structures and software queues. For tree data structures, Grove uses the same lowering machinery as Gyurgyik et al. [25], which supports user-specified memory layouts for ADT-defined accelerator trees [59]. For queues, Grove generates a struct-of-arrays (SoA), which enables vectorized/coalesced loads and stores to the queue, as noted in multiple wavefront renderers [35, 37, 52]. We do not implement the more complex queue layouts described by the MoonRay renderer [37], and instead match the layouts of the renderer we compare to in evaluation; the PBRT wavefront renderer uses SoA [52]. We leave exploration of memory layouts for the work queues, which could be co-designed with the schedule, as future work.

Vectorization. Vector instruction selection incorporates a number of peephole optimizations inspired by ispc [54] and Halide [55, 56]. For example, we use bounds analysis to lower integer division to floating point division when it is exact. We additionally support several optimizations for vectorization over tree data structures expressed in the layout language (e.g., transforming gathers on contiguous layouts into vector loads).

Queues: Appends, Size Inference, Double Buffering, and Allocation Hoisting. An append is an atomic add on the queue's counter that reserves an index, followed by a store of the state at that index. The drain iterates over $[0, \mathit{count})$ and resets the counter. Grove computes an upper bound on the size of a queue as the sum of the loop extents that can produce work for the queue. For single-producer queues (the common case in rendering workloads), this inference is simply the size of the single producer loop. For more complex patterns, e.g., multiple producers feeding the backedge to `t` in Figure 3, dataflow analysis is used to upper bound the size (e.g., `t` cannot grow larger than its previous size, as there is at most one unique path to insert into `t` from the drain of `t`). A queue that appends directly to itself (e.g., through a recursive call to the deferred function) is double-buffered, with a read-only and a write-only copy at every step. Queue allocations are

hoisted when possible: the queues in Listing 3 are allocated once outside the sequential loop over bands and reused across bands.

Hardware Binding. As in Halide, Grove launches parallel kernels via static closure capturing. Additionally, Grove targets NVIDIA's RT Cores [45] and OptiX [48] via the same code generation approach as Dr.Jit [28]. Texture units are targeted via the API provided by NVIDIA [47]. Extending Grove to support additional hardware models requires extending the hardware binding support via techniques similar to Mosaic [9] and Exo [26].

9 Evaluation

The primary goal of our evaluation is to show that Grove's programming model offers expressive power *without* sacrificing performance or portability. We show that schedules from hand-optimized parallel tree traversals [60, 77], path tracers [52], and Monte Carlo PDE solvers [61–63] can be concisely expressed in our programming model. Furthermore, we show that the performance of our generated code matches or exceeds hand-optimized CPU and GPU implementations. Lastly, we show that the optimal choice of schedule is *data-dependent*, which strengthens the justification for providing a scheduling language for productive design space exploration of the scheduling space.

9.1 Methodology

Hardware and software. All benchmarks were measured on an AMD Ryzen 9 9950X3D with 16 Zen 5 cores (32 threads with SMT), 48 KiB of L1 data cache and 1 MiB of L2 per core, 128 MiB of L3 across its two dies, and AVX-512/AVX2/SSE4 support, with 192 GB of DDR5, and an NVIDIA GeForce RTX 5090 (32 GB of GDDR7). The machine runs Ubuntu 24.04.4, and uses NVIDIA driver 615.71.09 (CUDA 13.4), CUDA toolkit 13.2 and OptiX 9.1. Grove compiles with LLVM 23.1.1, all generated code is compiled at `-O3` for the native target, and generated GPU code goes through LLVM's NVPTX backend. All benchmark drivers are compiled with Clang 22.1 at `-O3 -march=native -std=c++20`. We compile reference systems with their release configurations (`-O3 -march=native`): PBRT-v4 [52], Embree 4.4.1 [77], FCPW [60], and WoSX [61].[2] Fast math is disabled everywhere except for PBRT's GPU kernels, which PBRT builds with `--use_fast_math`; we match it on that comparison.

We match the reference system's tree exactly (topology and memory layout) to isolate the effects of the schedule for each benchmark. Section 11 discusses the interplay of layout and schedule.

Timing. Each measurement is the minimum of three runs after one warm-up. Times for all systems exclude loading the input and tree construction.

Verification. Every result is checked against the reference system's: traversal and solver outputs directly, rendered images by per-pixel difference.

Datasets. We evaluate on scenes and meshes from `pbrt-v4-scenes`[3], plus the six demo meshes of the WoSX repository [61] for the PDE solvers; all are summarized in Appendix C.

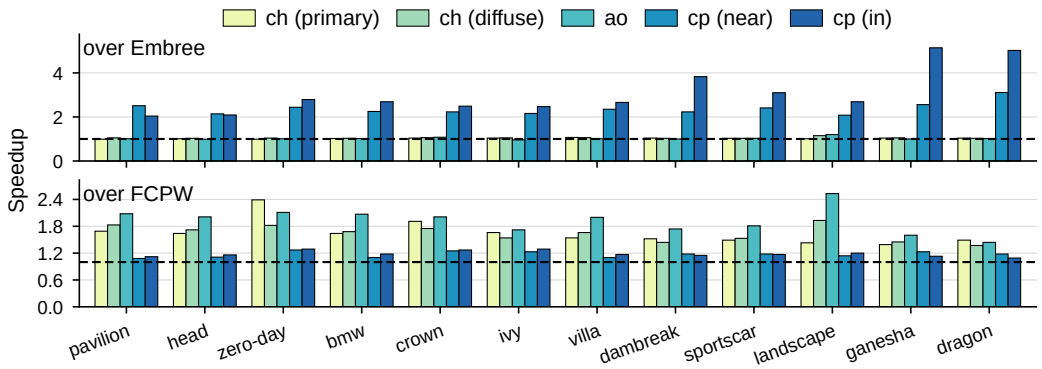
AI Use. We used Claude Fable 5.1 [5] for code development of some program transformations, schedule sweeps, and translation of the PBRT renderer into Grove's frontend. All code was reviewed by the authors. All results are verified to match up to floating point accumulation error.

9.2 High-Performance Tree Traversals

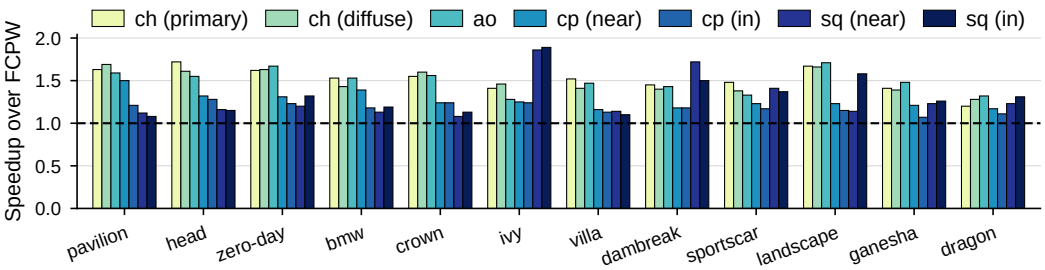
We compare Grove-generated traversals against two hand-optimized libraries: Embree [77] and Fastest Closest Points in the West (FCPW) [60]. We evaluate closest-hit (ch), any-hit (ao; ambient

[2] https://github.com/mmp/pbrt-v4, commit `eef3a6e`, https://github.com/RenderKit/embree, tag `v4.4.1`, https://github.com/rohan-sawhney/fcpw, commit `3b69ddb`, v1.2.0, https://github.com/nv-tlabs/wosx, commit `2716c27`

[3] https://github.com/mmp/pbrt-v4-scenes, commit `30cf4a0`



(a) CPU (1 core): Embree (8-wide BVH) and FCPW (4-wide). Geomean speedup over Embree: ch 1.02×/1.05× (primary/diffuse); ao 1.04×; cp 2.36×/2.95× (near/in). Over FCPW: ch 1.62×/1.63×; ao 1.90×; cp 1.17×/1.19×.



(b) GPU: FCPW (binary BVH). Geomean speedup: ch 1.53×/1.49× (primary/diffuse); ao 1.49×; cp 1.25×/1.18× (near/in); sq 1.27×/1.30× (near/in).

Fig. 8. Speedup of Grove-generated traversals over hand-optimized libraries with matching trees and schedules (higher is better, dashed line is parity). Queries: closest hit (ch) over primary and diffuse rays, any hit (ao) queries over shadow rays, and closest point (cp) and silhouette (sq), on points near or inside the mesh.

occlusion), closest-point (cp), and (GPU/FCPW only, not supported by Embree) silhouette (sq) queries on 12 meshes, with primary and diffuse rays for ray queries and points near or inside the mesh for point queries. Following the benchmarking of Aila and Laine [3], each mesh is traced from a 1024 × 1024 image of primary rays, and every primary hit casts one ambient-occlusion ray and one diffuse ray. Point queries (cp and sq) use points scattered uniformly over the bounding box, as FCPW's own test harness does [60], and points near the surface, as in the sampling scheme of Sawhney and Crane [62]. Figure 8 reports speedups over each library with matching schedules (Listing 6), so the differences are due to code generation alone.

CPU. Figure 8a compares against the fastest CPU tree each library provides: Embree's 8-wide and FCPW's 4-wide BVH, both with hand-vectorized traversals. Both libraries leave parallelism across queries to the application, so we measure one core to compare the quality of the vectorized traversals. On the queries each library was built for, Grove matches Embree on ray queries (ch and ao; worst case 0.95×, ao on ivy) and is modestly faster than FCPW on closest point queries. On the rest, Grove is faster on every mesh: up to 5.12× over Embree on cp (in) on ganesha, and up to 2.51× over FCPW on ao on landscape. We attribute the speedup over Embree to better vectorization of the loop over leaf data (`triangles::Leaf.data`) and the speedup over FCPW to better code generated for the vectorized sort of children (`trace.sort`). Listing 6 shows the ch schedule matching FCPW's

4-wide traversal. Embree's is similar, but it vectorizes over 8 children, prefetches in the loop over `triangles::Interior.cdrn`, and uses a larger stack depth.

Listing 6. Schedule matching FCPW's 4-wide closest-hit traversal, used for Figure 8a.

```
trace.sort(trace, |i:u32, r:Ray, bs:vec[AABB, 4]| dist(r, bs[i])) // visit in distance order
trace.vectorize(triangles::Interior.cdrn)    // vectorize child intersection and sort
trace.vectorize(triangles::Leaf.data)        // vectorize the loop over leaf data
trace.loopify(96)                            // recursion -> loop; stack depth 96 (FCPW's max)
```

GPU. FCPW's GPU backend supports only a binary BVH. We compare to it in Figure 8b. Grove is never slower than FCPW (minimum of 1.08× speedup across all queries). Geometric means of the speedups range from 1.18× on closest-point (in) to 1.53× on closest-hit (primary). Profiling attributes the speedup to `loopify`: the loop Grove generates caches the top of the traversal stack [20], reducing memory traffic relative to FCPW's hand-written loop.

9.3 Physically Based Rendering

PBRT is a research rendering system based on path tracing, an algorithm based on Monte Carlo simulation of a ray optics model of light [52]. This approach is the current state-of-the-art in computer graphics for image synthesis. We compare against PBRT [52]'s volumetric path tracer, which was implemented on both the GPU and the CPU, to provide evidence for three claims:

(1) Grove can express complex programs and complex schedules.
(2) Grove can outperform a hand-optimized renderer.
(3) The best choice of schedule is *scene-dependent*, in addition to machine-dependent.

Figure 13 in Appendix E illustrates the parallel dispatch graph of PBRT's GPU wavefront renderer, which executes a wavefront per pixel tile of the image. Grove's implementation of this PDG is 64 scheduling commands, where the GPU backend of PBRT took several months of development and approximately four thousand lines of code. Further, restructuring a significant amount of the C++ application code was required to run on both the CPU and the GPU [52]. PBRT is the largest program we evaluate, and while most transformations are whole-program transformations (and the Grove compiler itself is not parallelized), compiling the PBRT schedules takes only minutes, which is comparable to PBRT's sequential compile time (Figure 12 in Appendix D).

PBRT GPU Wavefront Comparison. Figure 9 shows the speedup of Grove-generated code over PBRT's matching renderer, across 30 scenes and 3 different samples per pixel parameters (spp roughly corresponds to work-per-pixel, which also improves image quality). Grove is faster across

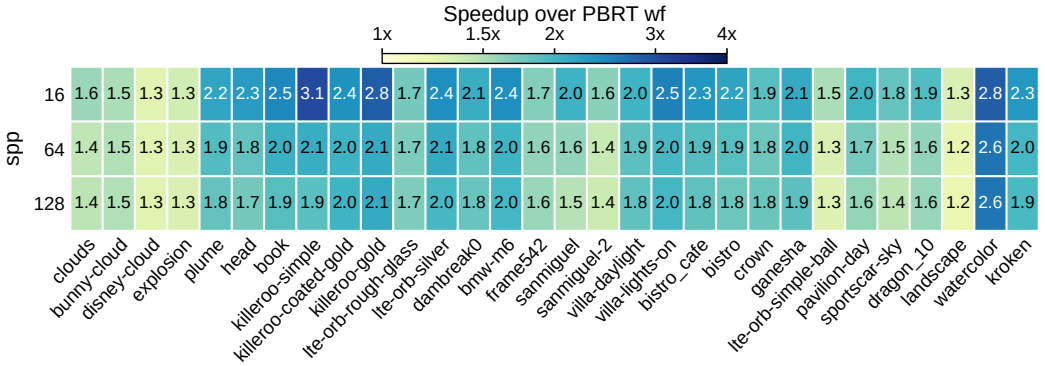
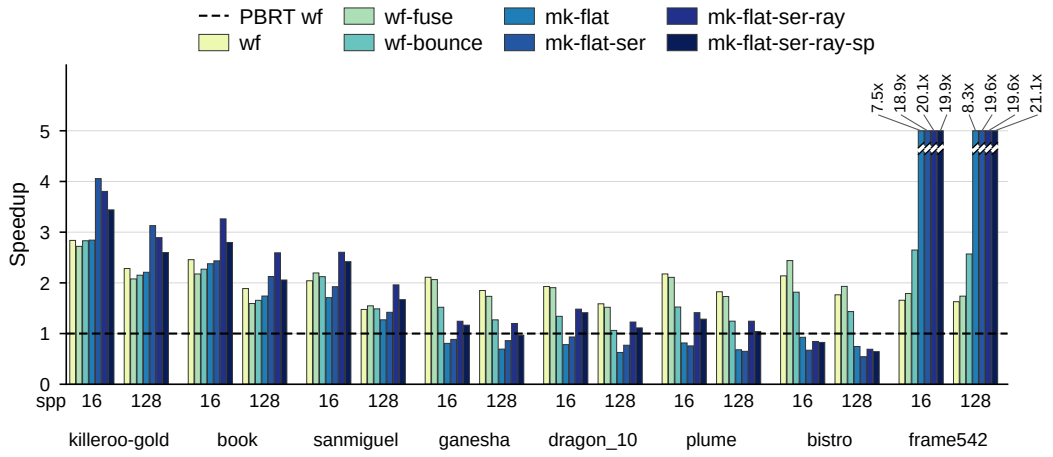
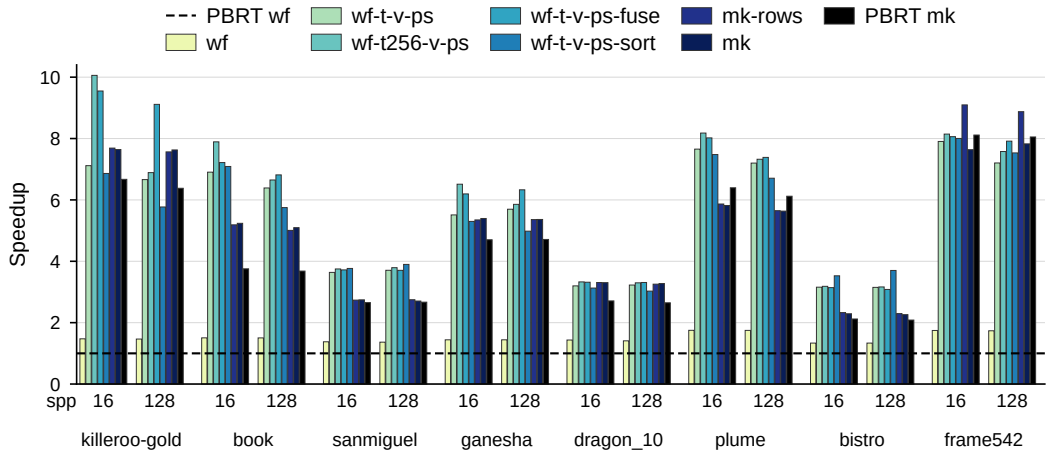


Fig. 9. Speedup of Grove-generated code over PBRT's GPU wavefront with a matching schedule.



(a) GPU schedule exploration: speedup over PBRT's wavefront. `wf` is Grove's matching schedule to PBRT. The best schedule is 1.56× faster than Grove's `wf` (geomean).



(b) CPU schedule exploration (32 threads): speedup over PBRT's wavefront. `wf` and `mk` match PBRT's wavefront and megakernel (black bar). The best schedule is 1.42× faster than PBRT's megakernel, 5.89× faster than PBRT's wavefront, and 3.93× faster than Grove's matching wavefront (geomeans).

Fig. 10. Schedule exploration across two hardware platforms. `wf` and `mk` match PBRT's wavefront and megakernel. The other schedules vary tiling (`-t`, `-t256`, `-rows`, and `-flat`), loop order (`-ps`, `-sp`), vectorization (`-v`), deferral (`-fuse`, `-bounce`), and sorting (`-sort`, `-ser`, and `-ray`).

all configurations. Profiling reveals the speedup to be a memory bandwidth win: Grove saves smaller continuations than the PBRT authors manually wrote, so queues are smaller and less data is loaded and stored across each otherwise-equal rendering stage (Table 3 in Appendix F).

PBRT Schedule Exploration. We now turn to exploration of schedules for PBRT's renderer. Figure 10a shows the speedup of 7 different schedules over PBRT's wavefront-only GPU renderer on 8 scenes chosen to span the dataset in geometry, materials, lights, and path depth (Table 2), at two samples-per-pixel (spp) rates. All schedules (including PBRT's) map the ray trace call to the GPU's RT Cores, and `wf` is our match of PBRT's schedule (Figure 9). The schedules we explored

include removal of different deferrals (e.g., `wf-fuse` fuses material evaluation into the trace kernel), and various megakernel (totally fused) variations, with and without shader execution reordering (SER) [46]. The PBRT-inspired schedule is not the strict winner: `wf` is the best schedule on three scenes, an SER megakernel on four, and `wf-fuse` on `bistro`. The best schedule is 1.56× faster than `wf` by geomean. On `frame542`, the SER megakernels are 12× and 13× faster than `wf` at 16 and 128 spp (20× and 21× faster than PBRT). Profiling shows that a very small number of incredibly long running paths per pixel tile leaves every wavefront tile long-running but low-occupancy; the megakernel launches all paths together, and SER compacts the long-running threads.

Figure 10b shows the speedup of 7 CPU schedules versus PBRT's CPU wavefront renderer on the same 8 scenes and sample rates. PBRT also offers a megakernel renderer on the CPU (black bar), which significantly outperforms PBRT's CPU wavefront. Our matching schedules (`wf` and `mk`) reproduce both, and `mk` is 1.10× faster than PBRT's megakernel by geomean. The best schedule is 1.42× faster than PBRT's megakernel by geomean (1.10–2.10×) and 1.29× faster than `mk`. On seven of the eight scenes the best schedule is a per-thread vectorized wavefront inspired by the MoonRay renderer [37], which tiles the image into small tiles such that work queues fit into L2 cache, executes tiles in parallel on CPU cores, and uses queues to group work into vector lanes worth of work. The fastest variant depends on the scene and sample rate: small tiles (`wf-t256-v-ps`) at 16 spp, fused material evaluation (`wf-t-v-ps-fuse`) at 128 spp, and sorting (`wf-t-v-ps-sort`) on the two scenes with the most materials.

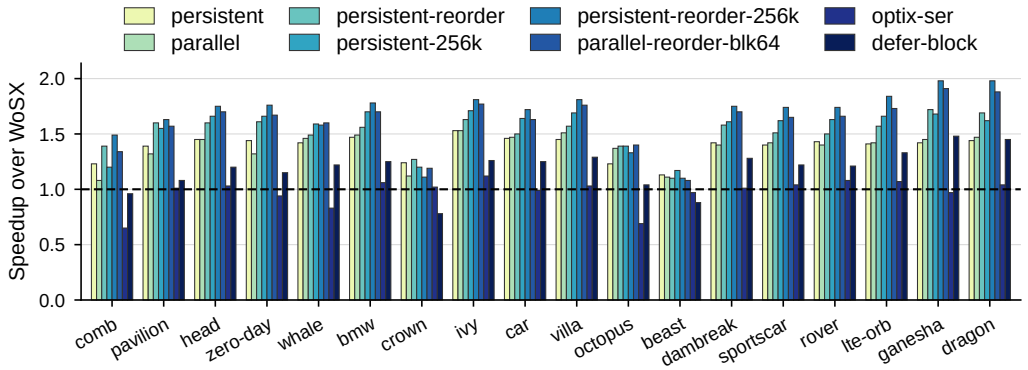
9.4 Monte Carlo PDE Solving

Walk on spheres (WoS) [62] and walk on stars (WoSt) [63] are Monte Carlo PDE solvers that, like a path tracer, estimate the solution at each point by averaging many independent random walks, each of which takes a data-dependent number of steps that perform spatial queries and evaluate the PDE's boundary functions. We compare against WoSX [61], the GPU implementation of both solvers. Unlike PBRT's wavefront GPU renderer, WoSX is a persistent-thread kernel [3]: it launches enough threads to fill the machine, which each claim a walk from an atomic counter in an effort to load balance walks across threads. This schedule is four scheduling commands in Grove (Listing 8).

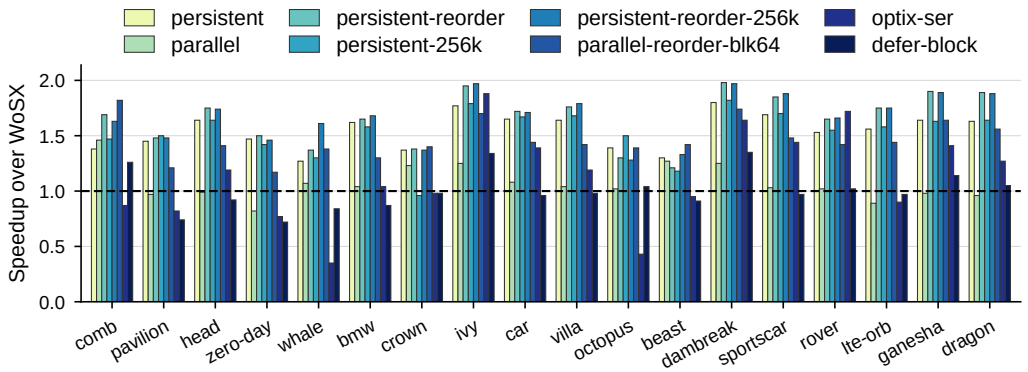
WoSX Comparison. Figure 11 shows the speedup of Grove-generated code over WoSX across 18 meshes. With the matching schedule (`persistent`), Grove is faster on every mesh: 1.38× by geomean on WoS and 1.54× on WoSt. Profiling attributes the win to occupancy. WoSX compiles a single kernel that supports every run-time option of its solver (e.g., also estimating gradients); on these benchmarks 70% of its instructions never execute, yet the GPU compiler still reserves registers for them, so WoSX needs 82 registers per thread. Grove's kernel, specialized to the program, needs 54 (the code WoSX actually executes needs the same), which raises achieved occupancy from 32% to 47% and hides more memory latency.

WoSX Schedule Exploration. The other schedules in Figure 11 vary the schedule (Appendix G), and as with rendering, the best schedule is data-dependent. On WoS, persistent threads make little difference: a plain launch of one thread per walk (`parallel`) is within 1% of `persistent`. The best schedule, `persistent-reorder-256k`, instead has each warp claim the walks of one point, which start from the same position, and launches twice as many threads; it is fastest on 14 of 18 meshes and 1.64× faster than WoSX by geomean. On WoSt, whose walks are longer and more varied in length, persistent threads make a large difference (`parallel` is 0.69× as fast as `persistent`), and five different schedules are the fastest on at least one mesh.

We also explored wavefront schedules that defer each step of a walk into a queue; Figure 11 shows only the best, `defer-block`, which is 0.85× as fast as `persistent` on WoS and 0.64× on WoSt. Wavefront execution performs poorly on this benchmark because the boundary functions of our test PDE are simple: a step is a few tree traversals, which diverge no matter which walks share



(a) Walk on spheres [62]. Geomean speedup: (matching) 1.38×; (best) 1.64× (up to 1.98×).



(b) Walk on stars [63]. Geomean speedup: (matching) 1.54×; (best) 1.66× (up to 1.98×).

Fig. 11. Speedup of Grove-generated GPU solvers over WoSX [61] on 18 inputs (higher is better, dashed line is parity). `persistent` matches WoSX's kernel; other schedules make changes explained in Appendix G.

a warp, and little else, so the divergence a deferral removes is not expensive enough to warrant the memory traffic of saving and restoring the step. However, a PDE with complex (divergent) boundary functions could justify wavefront execution, as materials do for PBRT.

10 Related Work

Divergence mitigation in renderers and traversals. Many of the transformations Grove exposes as scheduling directives were first described in manual restructurings of a particular renderer or traversal: wavefront path tracing [35], persistent threads [3, 24], ray reordering [40, 53], and, for memory coherence in traversals, packet tracing [76] and treelets [2]. Production renderers and traversal libraries [16, 37, 52, 60, 77] compose several of these with tiling and vectorization, but each composition is written into the application and tied to one machine.

Compilers for divergent graphics workloads. The GPU coroutines of Luisa [81] are the closest prior work to `defer`: suspension points annotated in a kernel are compiled into continuation frames, which one of a fixed set of schedulers (wavefront, persistent threads) executes. Dr.Jit [28] traces an SPMD program and executes the trace either fused into a megakernel or split at loops and virtual calls. Luisa's suspension points are baked into the kernel, so exploring a schedule means

editing the program, and Dr.Jit's split is a mode of the JIT; neither is a decision about a particular call, and neither compiler exposes tiling, vectorization, or specialization to compose with deferral. Rodent [50] and RaTrace [51] obtain portable renderers and traversals by partially evaluating [38] a library of higher-order traversal and shading variants against a scene and target, so their execution strategies devolve to a choice of library functions, not a composable scheduling space; we do think their partial evaluation would be useful in Grove for further data-specific specialization.

Stage-graph runtimes, stream languages, and parallel IRs. GRAMPS [72], Piko [49], and the GPU task runtimes of Tzeng et al. [74], Softshell [69], and Whippletree [70] schedule graphs of stages connected by queues with dynamic rates; GRAMPS inspired our treatment of data flow, and Piko goes furthest in separating a stage's logic from its scheduling (binning and ordering). A StreamIt [73] program is the same kind of graph, which its compiler [23] merges and splits to fit the machine, but a highly static one: stages push a fixed number of items per run, so queue expansion rates and loop counts are known at compile time; Brook [13], the predecessor to CUDA [39], maps kernels over fixed-length streams and has no queues at all. HPVM [34], a hierarchical dataflow graph with data-parallel nodes that targets CPUs, vector units, and GPUs, is the closest IR to a PDG, but is static in the same way: each edge carries one value per node instance, so a node cannot push to a queue on some instances and not others, and a drain cannot feed a queue upstream of itself. In each of these systems, the programmer writes the stages and the enqueues; a PDG derives both from the control flow of an SPMD program, so moving a stage boundary is a change to the schedule. Other parallel IRs target fork-join parallelism [57, 64, 67]; to the best of our knowledge, PDGs are the first parallel IR that models dynamic data flow alongside control flow and parallelism in one representation, and supports explicit transformation via a scheduling language.

User-schedulable languages. Grove's loop transformations are the standard loop transformations [8] of array DSLs [15, 26, 27, 55, 56], which schedule dense array programs whose loop structure is known at compile time, and whose schedules are pull-based: each producer is computed at some level of its consumer's loop nest. Divergent programs instead discover work at runtime, so Grove's schedules are *push*-based, driven by the producers whose work is deferred, and its new transformations rewrite control flow and data flow in addition to loops. Sparse tensor and graph compilers [31, 65, 71, 79, 80] schedule *irregular* programs, and some of their directives similarly target load imbalance and memory divergence, but the irregularity they schedule is in the *data*: index structures determine the iteration space, and a schedule transforms that iteration space. The divergence Grove is designed for is also in the *control*: every lane decides at runtime what execution path to take (e.g., how deeply to recurse) as a product of computation on the data, so its schedules must be dynamic and reactive.

SPMD compilation and vectorization. ispc [54], whole-function vectorization [29], and partial linearization [43] map SPMD code onto SIMD units as SIMT hardware does [39]: one instance per lane, with control flow linearized under masks. Our vectorization transformation builds directly on this work, including its notions of uniform and varying data and control flow. However, these techniques are standalone, leaving programmers to compose vectorization with other scheduling transformations by hand, as the MoonRay [37] developers did.

Machine models and portability. Sequoia [21] separates an algorithm from its mapping onto an explicit hierarchy of memories, Exo [26] moves hardware definitions into user libraries, and Mosaic [9] supports accelerator binding intermixed with loop scheduling for sparse tensor algebra; our hardware models (Section 6.3) follow each of these, including support for graphics accelerators.

Continuations. `defer` is a selective CPS conversion [7] followed by a defunctionalization [58] specialized per producer [10]. The transformation is not particularly novel, but we believe the user control and composition with loop scheduling and hardware bindings is.

11 Conclusion and Future Work

Parallel dispatch graphs represent control flow, data flow, and parallel loops in the same representation, and enable composable scheduling transformations to map SPMD programs onto SIMD hardware. The execution strategies of hand-optimized tree traversals, path tracers, and Monte Carlo PDE solvers are points in that scheduling space, and Grove matches or outperforms the hand-written originals. Just as Gyurgyik et al. [25] show that the optimal data layout is (machine, algorithm, data)-dependent, we show the same is true for the execution *schedule*. However, while our evaluation holds tree data layouts constant, we note that layouts and schedules are necessarily interdependent: vectorized tree traversals or queue drains work best when the corresponding data is stored in SoA, and some renderers argue for more complex queue layouts [37]. Together, these observations expose an interesting direction for future work: profile-guided co-optimization of layouts and schedules, across both the trees being traversed and the intermediate work queues that wavefront schedules drain. More generally, we believe that portable, high-performance code requires decoupling the program specification from the decisions that only affect performance (data structures, layouts, and schedules). Decoupling alone does not provide performance, however; it must be paired with a compiler that specializes the program to those decisions.

Acknowledgments

We thank Rohan Yadav for answering so, so, so many questions about GPUs and the visualization of parallel kernels. We also thank Atharva Chougule, James Dong, Olivia Hsu, Sang Yoon Kim, Rubens Lacouture, Katherine Mohr, Sai Gautham Ravipati, Dalia Shaaban, Nathan Sobotka, Rupanshu Soi, Shiv Sundram, Usman Tariq, and Bobby Yan for their helpful feedback on an early draft of this work. Additionally, we thank Sai Praveen Bangaru, Frédo Durand, Aaron Lefohn, Zander Majercik, Brennan Shacklett, and Venki Sivaram, for discussion on this work. Alexander and Christophe were supported by the Qualcomm Innovation Fellowship during part of this work. Alexander was also supported by the NVIDIA Graduate Fellowship and took part in this work while interning at NVIDIA under Matt and Rohan.

References

[1] Andrew Adams. 2021. Fast median filters using separable sorting networks. *ACM Trans. Graph.* 40, 4, Article 70 (July 2021), 11 pages. doi:10.1145/3450626.3459773

[2] Timo Aila and Tero Karras. 2010. Architecture considerations for tracing incoherent rays. In *Proceedings of the Conference on High Performance Graphics* (Saarbrucken, Germany) *(HPG '10)*. Eurographics Association, Goslar, DEU, 113–122.

[3] Timo Aila and Samuli Laine. 2009. Understanding the efficiency of ray traversal on GPUs. In *Proceedings of the Conference on High Performance Graphics 2009* (New Orleans, Louisiana) *(HPG '09)*. Association for Computing Machinery, New York, NY, USA, 145–149. doi:10.1145/1572769.1572792

[4] Frances E. Allen and John Cocke. 1972. A Catalogue of Optimizing Transformations. In *Design and Optimization of Compilers*, Randall Rustin (Ed.). Prentice-Hall, 1–30.

[5] Anthropic. 2026. *Claude Fable 5.1 & Claude Mythos 5.1 System Card.* Technical Report. Anthropic. https://www.anthropic.com/claude/fable

[6] Andrew W. Appel. 1998. SSA is functional programming. *SIGPLAN Not.* 33, 4 (April 1998), 17–20. doi:10.1145/278283.278285

[7] Andrew W. Appel. 2007. *Compiling with Continuations.* Cambridge University Press, USA.

[8] David F. Bacon, Susan L. Graham, and Oliver J. Sharp. 1994. Compiler transformations for high-performance computing. *ACM Comput. Surv.* 26, 4 (Dec. 1994), 345–420. doi:10.1145/197405.197406

[9] Manya Bansal, Olivia Hsu, Kunle Olukotun, and Fredrik Kjolstad. 2023. Mosaic: An Interoperable Compiler for Tensor Algebra. *Proc. ACM Program. Lang.* 7, PLDI, Article 122 (June 2023), 26 pages. doi:10.1145/3591236

[10] William Brandon, Benjamin Driscoll, Frank Dai, Wilson Berkow, and Mae Milano. 2023. Better Defunctionalization through Lambda Set Specialization. *Proc. ACM Program. Lang.* 7, PLDI, Article 146 (June 2023), 24 pages. doi:10.1145/3591260

[11] Matthias Braun, Sebastian Buchwald, Sebastian Hack, Roland Leißa, Christoph Mallon, and Andreas Zwinkau. 2013. Simple and efficient construction of static single assignment form. In *Proceedings of the 22nd International Conference on Compiler Construction* (Rome, Italy) *(CC'13)*. Springer-Verlag, Berlin, Heidelberg, 102–122. doi:10.1007/978-3-642-37051-9_6

[12] Preston Briggs, Keith D. Cooper, and Linda Torczon. 1992. Rematerialization. *SIGPLAN Not.* 27, 7 (July 1992), 311–321. doi:10.1145/143103.143143

[13] Ian Buck, Theresa Foley, Daniel Horn, Jeremy Sugerman, Kayvon Fatahalian, Mike Houston, and Pat Hanrahan. 2004. Brook for GPUs: stream computing on graphics hardware. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 777–786. doi:10.1145/1015706.1015800

[14] CCCL Development Team. 2023. *CCCL: CUDA C++ Core Libraries*. https://github.com/NVIDIA/cccl

[15] Chun Chen, Jacqueline Chame, and Mary Hall. 2008. *CHiLL: A framework for composing high-level loop transformations*. Technical Report. Citeseer.

[16] Per Christensen, Julian Fong, Tom Nettleship, Meghana Seshadri, Susan Salituro, Charlie Kilpatrick, Francisco Gonzalez, Srinath Ravichandran, Akshay Shah, Ethan Jaszewski, Stephen Friedman, James Burgess, and Trina M. Roy. 2025. RenderMan XPU: A Hybrid CPU+GPU Renderer for Interactive and Final-frame Rendering. *Computer Graphics Forum* (2025). doi:10.1111/cgf.70218

[17] F. Darema, D.A. George, V.A. Norton, and G.F. Pfister. 1988. A single-program-multiple-data computational model for EPEX/FORTRAN. *Parallel Comput.* 7, 1 (1988), 11–24. doi:10.1016/0167-8191(88)90094-4

[18] Zachary DeVito, Niels Joubert, Francisco Palacios, Stephen Oakley, Montserrat Medina, Mike Barrientos, Erich Elsen, Frank Ham, Alex Aiken, Karthik Duraisamy, Eric Darve, Juan Alonso, and Pat Hanrahan. 2011. Liszt: a domain specific language for building portable mesh-based PDE solvers. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis* (Seattle, Washington) *(SC '11)*. Association for Computing Machinery, New York, NY, USA, Article 9, 12 pages. doi:10.1145/2063384.2063396

[19] James Dong and Fredrik Kjolstad. 2026. A Compiler for Fused Relational Operations on Multisets. *Proc. ACM Program. Lang.* 10, PLDI, Article 205 (June 2026), 24 pages. doi:10.1145/3808283

[20] M. Anton Ertl. 1995. Stack caching for interpreters. In *Proceedings of the ACM SIGPLAN 1995 Conference on Programming Language Design and Implementation* (La Jolla, California, USA) *(PLDI '95)*. Association for Computing Machinery, New York, NY, USA, 315–327. doi:10.1145/207110.207165

[21] Kayvon Fatahalian, Daniel Reiter Horn, Timothy J. Knight, Larkhoon Leem, Mike Houston, Ji Young Park, Mattan Erez, Manman Ren, Alex Aiken, William J. Dally, and Pat Hanrahan. 2006. Sequoia: programming the memory hierarchy. In *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing* (Tampa, Florida) *(SC '06)*. Association for Computing Machinery, New York, NY, USA, 83–es. doi:10.1145/1188455.1188543

[22] Wilson W.L. Fung, Ivan Sham, George Yuan, and Tor M. Aamodt. 2007. Dynamic Warp Formation and Scheduling for Efficient GPU Control Flow. In *40th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2007)*. 407–420. doi:10.1109/MICRO.2007.30

[23] Michael I. Gordon, William Thies, Michal Karczmarek, Jasper Lin, Ali S. Meli, Andrew A. Lamb, Chris Leger, Jeremy Wong, Henry Hoffmann, David Maze, and Saman Amarasinghe. 2002. A stream compiler for communication-exposed architectures. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems* (San Jose, California) *(ASPLOS X)*. Association for Computing Machinery, New York, NY, USA, 291–303. doi:10.1145/605397.605428

[24] Kshitij Gupta, Jeff A. Stuart, and John D. Owens. 2012. A study of Persistent Threads style GPU programming for GPGPU workloads. In *2012 Innovative Parallel Computing (InPar)*. 1–14. doi:10.1109/InPar.2012.6339596

[25] Christophe Gyurgyik, Alexander J Root, and Fredrik Kjolstad. 2026. Decoupling Data Layouts from Bounding Volume Hierarchies. *Proceedings of the ACM on Programming Languages* 10, PLDI, Article 175 (June 2026). doi:10.1145/3808253

[26] Yuka Ikarashi, Gilbert Louis Bernstein, Alex Reinking, Hasan Genc, and Jonathan Ragan-Kelley. 2022. Exocompilation for productive programming of hardware accelerators. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (San Diego, CA, USA) *(PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 703–718. doi:10.1145/3519939.3523446

[27] Yuka Ikarashi, Kevin Qian, Samir Droubi, Alex Reinking, Gilbert Louis Bernstein, and Jonathan Ragan-Kelley. 2025. Exo 2: Growing a Scheduling Language. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1* (Rotterdam, Netherlands) *(ASPLOS '25)*. Association for Computing Machinery, New York, NY, USA, 426–444. doi:10.1145/3669940.3707218

[28] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, and Delio Vicini. 2022. DR.JIT: a just-in-time compiler for differentiable rendering. *ACM Trans. Graph.* 41, 4, Article 124 (July 2022), 19 pages. doi:10.1145/3528223.3530099

[29] Ralf Karrenberg and Sebastian Hack. 2011. Whole-function vectorization. In *International Symposium on Code Generation and Optimization (CGO 2011)*. 141–150. doi:10.1109/CGO.2011.5764682

[30] Gary A. Kildall. 1973. A unified approach to global program optimization. In *Proceedings of the 1st Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages* (Boston, Massachusetts) *(POPL '73)*. Association for Computing Machinery, New York, NY, USA, 194–206. doi:10.1145/512927.512945

[31] Fredrik Kjolstad, Shoaib Kamil, Stephen Chou, David Lugato, and Saman Amarasinghe. 2017. The tensor algebra compiler. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 77 (Oct. 2017), 29 pages. doi:10.1145/3133901

[32] Jens Knoop, Oliver Rüthing, and Bernhard Steffen. 1992. Lazy code motion. In *Proceedings of the ACM SIGPLAN 1992 Conference on Programming Language Design and Implementation* (San Francisco, California, USA) *(PLDI '92)*. Association for Computing Machinery, New York, NY, USA, 224–234. doi:10.1145/143095.143136

[33] Jens Knoop, Oliver Rüthing, and Bernhard Steffen. 1994. Partial dead code elimination. In *Proceedings of the ACM SIGPLAN 1994 Conference on Programming Language Design and Implementation* (Orlando, Florida, USA) *(PLDI '94)*. Association for Computing Machinery, New York, NY, USA, 147–158. doi:10.1145/178243.178256

[34] Maria Kotsifakou, Prakalp Srivastava, Matthew D. Sinclair, Rakesh Komuravelli, Vikram Adve, and Sarita Adve. 2018. HPVM: heterogeneous parallel virtual machine. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Vienna, Austria) *(PPoPP '18)*. Association for Computing Machinery, New York, NY, USA, 68–80. doi:10.1145/3178487.3178493

[35] Samuli Laine, Tero Karras, and Timo Aila. 2013. Megakernels considered harmful: wavefront path tracing on GPUs. In *Proceedings of the 5th High-Performance Graphics Conference* (Anaheim, California) *(HPG '13)*. Association for Computing Machinery, New York, NY, USA, 137–143. doi:10.1145/2492045.2492060

[36] Chris Lattner and Vikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-Directed and Runtime Optimization* (Palo Alto, California) *(CGO '04)*. IEEE Computer Society, USA, 75.

[37] Mark Lee, Brian Green, Feng Xie, and Eric Tabellion. 2017. Vectorized production path tracing. In *Proceedings of High Performance Graphics* (Los Angeles, California) *(HPG '17)*. Association for Computing Machinery, New York, NY, USA, Article 10, 11 pages. doi:10.1145/3105762.3105768

[38] Roland Leißa, Klaas Boesche, Sebastian Hack, Arsène Pérard-Gayot, Richard Membarth, Philipp Slusallek, André Müller, and Bertil Schmidt. 2018. AnyDSL: a partial evaluation framework for programming high-performance libraries. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 119 (Oct. 2018), 30 pages. doi:10.1145/3276489

[39] Erik Lindholm, John Nickolls, Stuart Oberman, and John Montrym. 2008. NVIDIA Tesla: A Unified Graphics and Computing Architecture. *IEEE Micro* 28, 2 (2008), 39–55. doi:10.1109/MM.2008.31

[40] Daniel Meister, Jakub Boksansky, Michael Guthe, and Jiri Bittner. 2020. On Ray Reordering Techniques for Faster GPU Ray Tracing. In *Symposium on Interactive 3D Graphics and Games* (San Francisco, CA, USA) *(I3D '20)*. Association for Computing Machinery, New York, NY, USA, Article 13, 9 pages. doi:10.1145/3384382.3384534

[41] Daniel Meister, Paritosh Kulkarni, Aaryaman Vasishta, and Takahiro Harada. 2024. HIPRT: A Ray Tracing Framework in HIP. *Proc. ACM Comput. Graph. Interact. Tech.* 7, 3, Article 44 (Aug. 2024), 18 pages. doi:10.1145/3675378

[42] Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. 2024. Walkin' Robin: Walk on Stars with Robin Boundary Conditions. *ACM Trans. Graph.* 43, 4, Article 41 (July 2024), 18 pages. doi:10.1145/3658153

[43] Simon Moll and Sebastian Hack. 2018. Partial control-flow linearization. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Philadelphia, PA, USA) *(PLDI 2018)*. Association for Computing Machinery, New York, NY, USA, 543–556. doi:10.1145/3192366.3192413

[44] Frank Mueller and David B. Whalley. 1995. Avoiding conditional branches by code replication. In *Proceedings of the ACM SIGPLAN 1995 Conference on Programming Language Design and Implementation* (La Jolla, California, USA) *(PLDI '95)*. Association for Computing Machinery, New York, NY, USA, 56–66. doi:10.1145/207110.207116

[45] NVIDIA Corporation. 2018. *NVIDIA Turing GPU Architecture: Graphics Reinvented.* Whitepaper WP-09183-001_v01. NVIDIA Corporation. https://images.nvidia.com/aem-dam/Solutions/design-visualization/technologies/turing-architecture/NVIDIA-Turing-Architecture-Whitepaper.pdf

[46] NVIDIA Corporation. 2022. *Shader Execution Reordering.* https://d29g4g2dyqv443.cloudfront.net/sites/default/files/akamai/gameworks/ser-whitepaper.pdf Version 1.0.

[47] NVIDIA Corporation. 2025. CUDA C++ Programming Guide, Section 19. Texture Fetching. NVIDIA Developer Documentation. https://docs.nvidia.com/cuda/cuda-c-programming-guide/#texture-fetching

[48] Steven G. Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, and Martin Stich. 2010. OptiX: a general purpose ray tracing engine. *ACM Trans. Graph.* 29, 4, Article 66 (July 2010), 13 pages. doi:10.1145/1778765.1778803

[49] Anjul Patney, Stanley Tzeng, Kerry A. Seitz, and John D. Owens. 2015. Piko: a framework for authoring programmable graphics pipelines. *ACM Trans. Graph.* 34, 4, Article 147 (July 2015), 13 pages. doi:10.1145/2766973

[50] Arsène Pérard-Gayot, Richard Membarth, Roland Leißa, Sebastian Hack, and Philipp Slusallek. 2019. Rodent: generating renderers without writing a generator. *ACM Trans. Graph.* 38, 4, Article 40 (July 2019), 12 pages. doi:10.1145/3306346.3322955

[51] Arsène Pérard-Gayot, Martin Weier, Richard Membarth, Philipp Slusallek, Roland Leißa, and Sebastian Hack. 2017. RaTrace: simple and efficient abstractions for BVH ray traversal algorithms. *SIGPLAN Not.* 52, 12 (Oct. 2017), 157–168. doi:10.1145/3170492.3136044

[52] Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically Based Rendering: From Theory to Implementation* (4 ed.). MIT Press. https://mitpress.mit.edu/9780262048026/physically-based-rendering/

[53] Matt Pharr, Craig Kolb, Reid Gershbein, and Pat Hanrahan. 1997. Rendering complex scenes with memory-coherent ray tracing. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 101–108. doi:10.1145/258734.258791

[54] Matt Pharr and William R. Mark. 2012. ispc: A SPMD compiler for high-performance CPU programming. In *2012 Innovative Parallel Computing (InPar)*. 1–13. doi:10.1109/InPar.2012.6339601

[55] Jonathan Ragan-Kelley, Andrew Adams, Sylvain Paris, Marc Levoy, Saman Amarasinghe, and Frédo Durand. 2012. Decoupling algorithms from schedules for easy optimization of image processing pipelines. *ACM Trans. Graph.* 31, 4, Article 32 (July 2012), 12 pages. doi:10.1145/2185520.2185528

[56] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Seattle, Washington, USA) *(PLDI '13)*. Association for Computing Machinery, New York, NY, USA, 519–530. doi:10.1145/2491956.2462176

[57] Mike Rainey, Ryan R. Newton, Kyle Hale, Nikos Hardavellas, Simone Campanoni, Peter Dinda, and Umut A. Acar. 2021. Task parallel assembly language for uncompromising parallelism. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (Virtual, Canada) *(PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 1064–1079. doi:10.1145/3453483.3460969

[58] John C. Reynolds. 1972. Definitional interpreters for higher-order programming languages. In *Proceedings of the ACM Annual Conference - Volume 2* (Boston, Massachusetts, USA) *(ACM '72)*. Association for Computing Machinery, New York, NY, USA, 717–740. doi:10.1145/800194.805852

[59] Alexander J Root, Christophe Gyurgyik, Purvi Goel, Kayvon Fatahalian, Jonathan Ragan-Kelley, Andrew Adams, and Fredrik Kjolstad. 2026. Bonsai: Compiling Queries to Pruned Tree Traversals. *Proceedings of the ACM on Programming Languages* 10, PLDI, Article 178 (June 2026). doi:10.1145/3808256

[60] Rohan Sawhney. 2021. *FCPW: Fastest Closest Points in the West.* https://github.com/rohan-sawhney/fcpw

[61] Rohan Sawhney. 2026. *Walk on Spheres Extensions (WoSX).* https://github.com/nv-tlabs/wosx

[62] Rohan Sawhney and Keenan Crane. 2020. Monte Carlo geometry processing: a grid-free approach to PDE-based methods on volumetric domains. *ACM Trans. Graph.* 39, 4, Article 123 (Aug. 2020), 18 pages. doi:10.1145/3386569.3392374

[63] Rohan Sawhney, Bailey Miller, Ioannis Gkioulekas, and Keenan Crane. 2023. Walk on Stars: A Grid-Free Monte Carlo Method for PDEs with Neumann Boundary Conditions. *ACM Trans. Graph.* 42, 4, Article 80 (July 2023), 20 pages. doi:10.1145/3592398

[64] Tao B. Schardl, William S. Moses, and Charles E. Leiserson. 2019. Tapir: Embedding Recursive Fork-join Parallelism into LLVM's Intermediate Representation. *ACM Trans. Parallel Comput.* 6, 4, Article 19 (Dec. 2019), 33 pages. doi:10.1145/3365655

[65] Ryan Senanayake, Changwan Hong, Ziheng Wang, Amalee Wilson, Stephen Chou, Shoaib Kamil, Saman Amarasinghe, and Fredrik Kjolstad. 2020. A sparse iteration space transformation framework for sparse tensor algebra. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 158 (Nov. 2020), 30 pages. doi:10.1145/3428226

[66] Peter Shirley, Trevor David Black, and Steve Hollasch. 2025. Ray Tracing in One Weekend. Online book. https://raytracing.github.io/books/RayTracingInOneWeekend.html

[67] Harini Srinivasan, James Hook, and Michael Wolfe. 1993. Static single assignment for explicitly parallel programs. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Charleston, South Carolina, USA) *(POPL '93)*. Association for Computing Machinery, New York, NY, USA, 260–272. doi:10.1145/158511.158644

[68] Guy Lewis Steele. 1977. Debunking the "expensive procedure call" myth or, procedure call implementations considered harmful or, LAMBDA: The Ultimate GOTO. In *Proceedings of the 1977 Annual Conference* (Seattle, Washington) *(ACM '77)*. Association for Computing Machinery, New York, NY, USA, 153–162. doi:10.1145/800179.810196

[69] Markus Steinberger, Bernhard Kainz, Bernhard Kerbl, Stefan Hauswiesner, Michael Kenzel, and Dieter Schmalstieg. 2012. Softshell: dynamic scheduling on GPUs. *ACM Trans. Graph.* 31, 6, Article 161 (Nov. 2012), 11 pages. doi:10.1145/2366145.2366180

[70] Markus Steinberger, Michael Kenzel, Pedro Boechat, Bernhard Kerbl, Mark Dokter, and Dieter Schmalstieg. 2014. Whippletree: task-based scheduling of dynamic workloads on the GPU. *ACM Trans. Graph.* 33, 6, Article 228 (Nov. 2014), 11 pages. doi:10.1145/2661229.2661250

[71] Michelle Mills Strout, Mary Hall, and Catherine Olschanowsky. 2018. The Sparse Polyhedral Framework: Composing Compiler-Generated Inspector-Executor Code. *Proc. IEEE* 106, 11 (2018), 1921–1934. doi:10.1109/JPROC.2018.2857721

[72] Jeremy Sugerman, Kayvon Fatahalian, Solomon Boulos, Kurt Akeley, and Pat Hanrahan. 2009. GRAMPS: A programming model for graphics pipelines. *ACM Trans. Graph.* 28, 1, Article 4 (Feb. 2009), 11 pages. doi:10.1145/1477926.1477930

[73] William Thies, Michal Karczmarek, and Saman P. Amarasinghe. 2002. StreamIt: A Language for Streaming Applications. In *Proceedings of the 11th International Conference on Compiler Construction (CC '02)*. Springer-Verlag, Berlin, Heidelberg, 179–196.

[74] Stanley Tzeng, Anjul Patney, and John D. Owens. 2010. Task management for irregular-parallel workloads on the GPU. In *Proceedings of the Conference on High Performance Graphics* (Saarbrucken, Germany) *(HPG '10)*. Eurographics Association, Goslar, DEU, 29–37.

[75] UXL Foundation. 2026. oneTBB `parallel_sort`. oneAPI Specification. https://oneapi-spec.uxlfoundation.org/specifications/oneapi/v1.4-rev-1/elements/onetbb/source/algorithms/functions/parallel_sort_func Accessed November 2026.

[76] Ingo Wald, Philipp Slusallek, Carsten Benthin, and Markus Wagner. 2001. Interactive Rendering with Coherent Ray Tracing. *Computer Graphics Forum* 20, 3 (2001), 153–165. doi:10.1111/1467-8659.00508

[77] Ingo Wald, Sven Woop, Carsten Benthin, Gregory S. Johnson, and Manfred Ernst. 2014. Embree: a kernel framework for efficient CPU ray tracing. *ACM Trans. Graph.* 33, 4, Article 143 (July 2014), 8 pages. doi:10.1145/2601097.2601199

[78] Mark N. Wegman and F. Kenneth Zadeck. 1991. Constant propagation with conditional branches. *ACM Trans. Program. Lang. Syst.* 13, 2 (April 1991), 181–210. doi:10.1145/103135.103136

[79] Yunming Zhang, Mengjiao Yang, Riyadh Baghdadi, Shoaib Kamil, Julian Shun, and Saman Amarasinghe. 2018. GraphIt: a high-performance graph DSL. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 121 (Oct. 2018), 30 pages. doi:10.1145/3276491

[80] Tuowen Zhao, Tobi Popoola, Mary Hall, Catherine Olschanowsky, and Michelle Strout. 2022. Polyhedral Specification and Code Generation of Sparse Tensor Contraction with Co-iteration. *ACM Trans. Archit. Code Optim.* 20, 1, Article 16 (Dec. 2022), 26 pages. doi:10.1145/3566054

[81] Shaokun Zheng, Xin Chen, Zhong Shi, Ling-Qi Yan, and Kun Xu. 2024. GPU Coroutines for Flexible Splitting and Scheduling of Rendering Tasks. *ACM Trans. Graph.* 43, 6, Article 281 (Nov. 2024), 24 pages. doi:10.1145/3687766

A RTIOW Renderer Intermediate Representation

Listing 7. Listing 1's path tracer lowered to Grove's IR. **ret**s are inlined and types are omitted for brevity.

```
func interact(ray, hit):
 m = hit.material
 cjmp m [lamb(ray, hit), metal(ray, hit),
         diel(ray, hit), light(hit)]
lamb(albedo, ray, hit):
 call rngu() lamb_c(albedo, hit)
lamb_c(u, albedo, hit):
 d = hit.n + u
 jmp scatter(d, albedo, hit)
metal(albedo, fuzz, ray, hit):
 call reflect(ray.dir, hit.n)
      metal_c(albedo, fuzz, hit)
metal_c(rd, albedo, fuzz, hit):
 call rngu() metal_cc(rd, albedo, fuzz, hit)
metal_cc(u, rd, albedo, fuzz, hit):
 fu = fuzz * u
 d = rd + fu
 call dot(d, hit.n) metal_ccc(albedo, hit, d)
metal_ccc(c, albedo, hit, d):
 g = c <= 0
 cjmp g [scatter(d, albedo, hit), // false
         metal_a()]               // true
metal_a(): // grazing
 a = Absorb{0}
 ret a
diel(eta, ray, hit):
 call ref(ray.dir, hit, eta) scatter(1, hit)
light(tex, hit):
 call eval(tex, hit.pt) light_c()
light_c(e):
 a = Absorb{e}
 ret a

// jmp target (lamb_c), cjmp arm (metal_ccc),
// and call continuation (diel)
scatter(d, albedo, hit):
 r = Ray{hit.pt, d}
 s = Scatter{r, albedo, 0}
 ret s
```

```
func Li(L, beta, ray, n, scene):
 c = n == 0
 cjmp c [li_tr(L, beta, ray, n, scene), ret L]
li_tr(L, beta, ray, n, scene):
 call trace(ray, scene)
      li_hit(L, beta, ray, n, scene)
li_hit(h, L, beta, ray, n, scene):
 cjmp h [li_miss(L, beta, ray),
         li_some(L, beta, ray, n, scene)]
li_miss(L, beta, ray):
 call background(ray) li_emit(L, beta)
li_some(hit, L, beta, ray, n, scene):
 call interact(ray, hit)
      li_resp(L, beta, n, scene)
li_resp(resp, L, beta, n, scene):
 cjmp resp [li_emit(L, beta),
            li_scatter(L, beta, n, scene)]
li_emit(e, L, beta): // miss and Absorb
 t = beta * e
 r = L + t
 ret r
li_scatter(dir, a, e, L, beta, n, scene):
 t = beta * e
 L1 = L + t
 beta1 = beta * a
 n1 = n - 1
 call Li(L1, beta1, dir, n1, scene) ret

func render(width, height, out, depth, scene):
 N = width * height
 par p [0, N, 1] body(out, depth, scene) ret
body(p, out, depth, scene):
 call camera_ray(p) cont(p, out, depth, scene)
cont(r, p, out, depth, scene):
 call Li(0.0, 1.0, r, depth, scene)
      body_store(p, out)
body_store(v, p, out):
 store out[p], v
 yield
```

B Defer Implementation

Algorithm 1 is the deferral transformation of Section 7.3; $Defer_q$ is the `Defer` arm for queue q, and PlaceAtOwner performs the legality check described in Section 7.3.

Algorithm 1 Defer Implementation, shown for a call edge `call g(a) k(b)`. For a call continuation the stored state is the return value `v` with `b`, and the drain resumes at `k(v, b)`; for a dispatch arm it is the arm's arguments, and the drain resumes at the arm. $func(q)$ is the function containing the loop that owns q; w is the stored state, which the drain D passes to every block (elided); `Defer_q` is the `Defer` arm for queue q.

```
1:  function DEFER(f, g, q)                          ▷ c = f.defer(g, q)
2:    for all s = call g(a) k(b) ∈ f do
3:      rewrite s to defer (s, a ++ b) q
4:      q.producers += s
5:    PROPAGATE(f, q)
6:    return EMITDRAIN(q)                            ▷ the cursor c

7:  function PROPAGATE(f, q)                         ▷ up to q's loop
8:    if f = func(q) then return                     ▷ loop body yields
9:    if callers(f) = ∅ then error                   ▷ g not under q's owner
10:   if Defer_q ∈ rettype(f) then return
11:   if rettype(f) ≠ Result then
12:     rettype(f) ← Result                          ▷ Done(T) | Defer_*(Index)
13:     rewrite each ret v to ret Done(v)
14:   rettype(f) += Defer_q(Index)                   ▷ one arm per queue
15:   for all s' = call f(a) k(b) ∈ callers(f) do
16:     if k ≠ ret then EMITPATCH(s', q)
17:     PROPAGATE(func(s'), q)

18: function EMITPATCH(s', q)                        ▷ s' = call f(a) k(b)
19:   assert func(s') ∉ callees*(f)                  ▷ non-tail recursion
20:   rewrite s' to call f(a) k'(b)
21:   emit k'(x, b):                                 ▷ x : Result
            cjmp x [k(b), p_q(b), ...]               ▷ Done→k, Defer_q→p_q
22:   emit p_q(i, b):                                ▷ save this frame
            q[i].frames += (s', b)
            ret Defer_q(i)                           ▷ yield in q's loop body
23:   q.patches += (s', k)

24: function EMITDRAIN(q)                            ▷ rebuilt per directive
25:   emit D(w):                                     ▷ an arm per s ∈ q.producers
            cjmp w.tag [L_s(w), ...]
26:   for all s = call g(a) k(b) ∈ q.producers do
27:     emit L_s(w): call g(w.a) k(w.b)              ▷ resume the call
28:     EMITCOPY(k)
29:   for all (s', k) ∈ q.patches do EMITCOPY(k)
30:   EMITRETURN(q)
31:   q.drain ← par w ∈ q do D(w)
32:   PLACEATOWNER(q)
33:   return D

34: function EMITCOPY(k)                             ▷ once per k
35:   blocks(D) ∪= copy(reach(k))                    ▷ k and its successors
36:   rewrite each ret x in the copy to jmp Return(x)

37: function EMITRETURN(q)                           ▷ returning inside D
38:   emit Return(x):                                ▷ x : Result
            cjmp x [Resume(), Forward_q'(), ...]     ▷ per queue q'
39:   emit Resume(v):                                ▷ innermost frame first
            (t, b) = pop_front(w.frames)
            cjmp t [k(v, b), ...]                    ▷ (s', k) ∈ q.patches
40:   emit Forward_q'(i):                           ▷ deferred again, into q'
            q'[i].frames ++= w.frames; yield

41: function PLACEATOWNER(q)                         ▷ legality, Section 7.3
42:   insert q.drain at owner(q) after every producer
43:   for all S between producer P and q.drain, d ∈ data(D) do
44:     assert access_S(d) = access_P(d)
45:   if D feeds a queue upstream then loop until empty
```

C Dataset Statistics

Table 1. Meshes of the tree traversal (Section 9.2) and Monte Carlo PDE (Section 9.4) benchmarks. Meshes come from pbrt-v4-scenes [52] and the six demo meshes of the WoSX repository [61] (†), used only for the PDE benchmarks. Figure 8 omits lte-orb and Figure 11 omits landscape, which are degenerate inputs for those benchmarks.

mesh	file	triangles
comb†	`electrostatics/data/mems-comb-electrodes.obj`	600
pavilion	`barcelona-pavilion/geometry/mesh_00099.ply`	11,366
head	`head/geometry/head.ply`	17,674
zero-day	`zero-day/geometry/9dbbf607. . . .ply`	48,960
whale†	`potential_flow/data/whale.obj`	99,648
bmw	`bmw-m6/geometry/mesh_00037.ply`	110,592
crown	`crown/geometry/mesh_00397.ply`	155,520
ivy	`sanmiguel/geometry/enredadera_00001.ply`	179,603
car†	`potential_flow/data/car.obj`	324,640
villa	`villa/geometry/mesh_00072.ply`	390,784
octopus†	`geometric_deformation/data/octopus.obj`	497,236
beast†	`geometric_deformation/data/beast.obj`	606,136
dambreak	`dambreak/geometry/dambreak1.ply`	1,015,024
sportscar	`sportscar/geometry/OSD_CarPoly_0000_m028.ply`	1,091,232
rover†	`thermal_conduction/data/rover.zip`	1,616,316
landscape	`landscape/geometry/mesh_00370.ply`	1,916,928
lte-orb	`lte-orb/geometry/mesh-3.ply.gz`	3,423,232
ganesha	`ganesha/geometry/ganesha.ply.gz`	4,323,658
dragon	`sssdragon/geometry/dragon.ply.gz`	7,219,045

Table 2. PBRT evaluation scenes. Scenes named by filename in pbrt-v4-scenes [52], except (for brevity) sanmiguel (sanmiguel/sanmiguel-courtyard), sanmiguel-2 (sanmiguel/sanmiguel-courtyard-second), bistro (bistro/bistro_vespa), landscape (landscape/view-0), watercolor (watercolor/camera-1), kroken (kroken/camera-1).

scene	triangles	meshes	instances	materials	textures	lights	resolution	depth
clouds	0	0	–	1	0	1	1920×1080	50
bunny-cloud	0	0	–	2	0	1	1920×1080	50
disney-cloud	0	0	–	2	0	2	1280×720	100
explosion	0	0	–	2	0	1	1300×1800	5
plume	2	1	–	3	0	1	1200×1200	5
head	17,674	1	–	2	3	1	1920×1080	2
book	21,246	3	–	4	9	2	1920×1080	5
killeroo-simple	66,532	4	–	5	0	1	700×700	5
killeroo-coated-gold	532,230	4	–	5	6	2	1368×1026	5
killeroo-gold	532,230	4	–	5	6	2	1368×1026	5
lte-orb-rough-glass	597,812	8	–	5	6	2	1200×1200	30
lte-orb-silver	597,812	8	–	5	6	2	1200×1200	5
dambreak0	620,956	8	–	3	0	1	1920×1080	64
bmw-m6	792,746	114	–	29	0	1	1400×1000	8
frame542	1,317,820	229	5	2	0	1	1920×840	256
sanmiguel	2,503,021	796	69,040	300	1,069	1	1600×1100	5
sanmiguel-2	2,503,044	808	159,924	300	1,069	1	1600×1100	5
villa-daylight	2,624,966	358	146	69	159	1	1800×870	10
villa-lights-on	2,624,966	358	146	69	159	8	1000×444	5
bistro_cafe	2,852,758	1,693	–	133	409	20,639	1920×1080	5
bistro	2,852,758	1,693	–	133	409	20,639	1920×1080	5
crown	3,540,226	798	–	52	23	12	1000×1400	100
ganesha	4,323,668	6	–	2	3	5	1440×1440	5
lte-orb-simple-ball	5,237,888	4	–	3	0	1	1200×1200	5
pavilion-day	5,311,719	120	43	39	127	1	1600×850	5
sportscar-sky	7,157,264	48	–	49	0	1	1920×1080	5
dragon_10	7,219,557	2	–	3	0	1	1366×1024	5
landscape	27,885,845	370	30,034	638	1,611	1	3840×2160	5
watercolor	58,383,675	944	54	261	412	2	1920×1440	15
kroken	287,904,585	248	10	80	164	2,176	1500×1500	15

D PBRT Compile Times

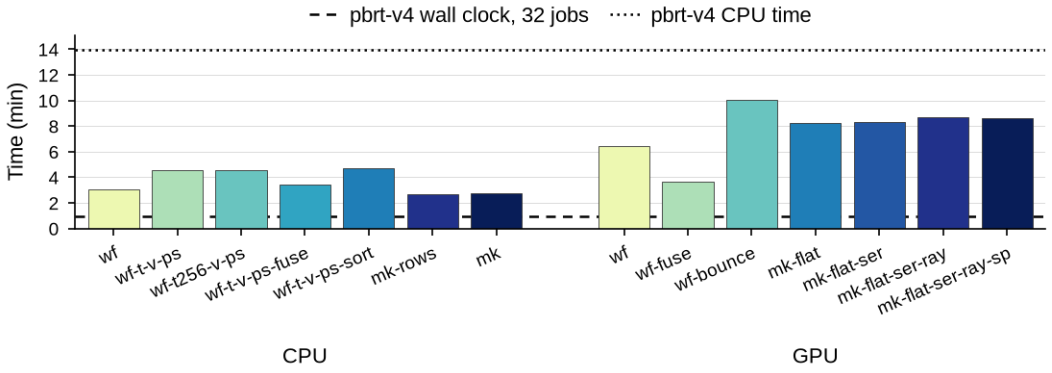


Fig. 12. Grove compile times for the schedules used in Section 9.3. Grove compilation is not parallelized, but PBRT is also compiling features (e.g., an interactive viewport) that our implementation does not include.

E PBRT's Parallel Dispatch Graph

Figure 13 shows PBRT's GPU wavefront integrator [52, Ch. 15.2] as the PDG that Grove's 64-directive schedule produces; the graph is the body of one tile (`band`, `i`) of the pixel loop. `camera_ray` fills the ray queue `r`; `ray_samples` draws each ray's samples, and the drain of `r` runs on OptiX threads [48] and traces on the RT cores [45]. Every edge out of the dispatches that follow is deferred: misses go to `e`; rays inside a medium to `s`, whose kernel scatters into `S` or passes to the surface hit; and hits dispatch on the material, emissive hits into `a`, interface materials back into `r`, and the remaining 22 materials into the queues $m_1 \ldots m_n$ of one `specialize`, so each `bsdf` kernel is compiled for one material. Subsurface materials defer again into `b`, whose probe rays `P` are traced on the RT cores and resumed by `sss`. Every kernel that extends a path appends to `r` while `d < D`, and every kernel that samples a light appends to the shadow queue `w`, also traced on the RT cores. `r` is the only queue its own drain can feed (the interface edge), so it is double-buffered, and the graph re-executes until `r` is empty (Section 3).

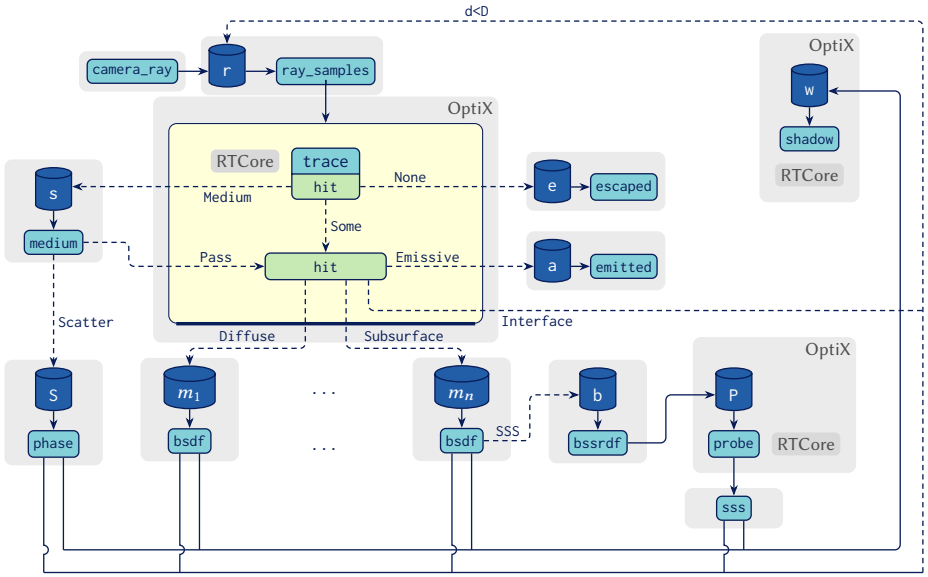


Fig. 13. Inner tile (with band and `i` nested loops) of PBRT's GPU path integrator [52, Ch. 15.2] expressed as a PDG. Surface hits are binned into per-material queues (two shown), so every `bsdf` kernel is specialized to a material. `r` is one logical queue that PBRT double-buffers, as it has a self-edge (its own drain can push to it).

F Queue Size Comparison

Table 3 compares the bytes per stored state in PBRT's wavefront integrator and in Grove's matching schedule. Grove's states are smaller for every queue but one, for two mechanical reasons: PBRT stores each boolean in 4 bytes, where Grove packs flags into the bits of one byte, and PBRT stores several 8-byte pointers (e.g., to a material) where Grove stores the 32-bit index of the same object in the per-material array. The exception is BSSRDF evaluation, where PBRT's GPU integrator stores an approximation of the BSSRDF state that does not match the semantics of its CPU integrator; Grove, which reproduces the CPU semantics, carries 12 more bytes.

Queue	PBRT	Grove
Rays	184	173
Medium samples	364	215
Medium scatters	120	112
Escaped rays	164	145
Emissive hits	192	169
Material hits (each of 22)	252	247
Shadow rays	124	96
BSSRDF evaluation	168	180
Subsurface probes	308	184
Subsurface exits	308	220
Exit shadow rays	124	96

Table 3. Bytes per stored state in PBRT's wavefront integrator and in Grove's matching schedule. PBRT reuses one queue for the subsurface probes and exits, and its shadow queue for both shadow traces.

G WoSX Schedules

Listing 8 is the schedule matching WoSX's persistent kernel. The solver has two parallel loops, over points (`i`) and over walks (`w`); the schedule reorders them so walks are outermost, collapses them into a single loop over (walk, point) work items so that consecutive items are consecutive points (WoSX's order), and makes that loop persistent. The tree traversals use the schedule of Section 9.2.

The other schedules of Figure 11 each change one or two scheduling commands. `parallel` drops `persistent` and launches one thread per work item. `persistent-reorder` collapses with points outermost instead, so that the lanes of a warp claim the walks of one point. `persistent-256k` uses twice as many persistent threads, and `persistent-reorder-256k` does both. The schedule `parallel-reorder-blk64` is a plain launch with points outermost and 64-thread blocks, so that inputs with few points still fill the GPU. `optix-ser` binds the walks to OptiX threads and sorts the walk's recursive call by the Morton code of its position, which lowers to shader execution reordering [46] before every step. `defer-block` binds points to GPU blocks and defers each step of a walk into a per-point queue in shared memory, which is drained in rounds.

```
solve_all.reorder(i, w)                  // walks outside points
k = solve_all.collapse(w, i)             // one work item per (walk, point)
t = solve_all.persistent(k, 131072)      // 131072 workers claim items from a counter
solve_all.gpu_tile(t, 256)               // 256 threads per block
```

Listing 8. Schedule matching WoSX's persistent GPU kernel (`persistent` in Figure 11).