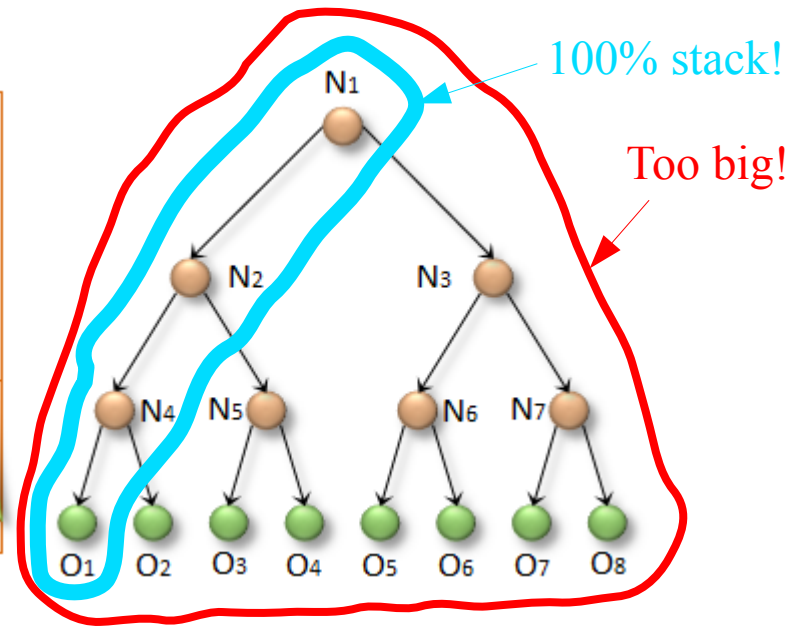
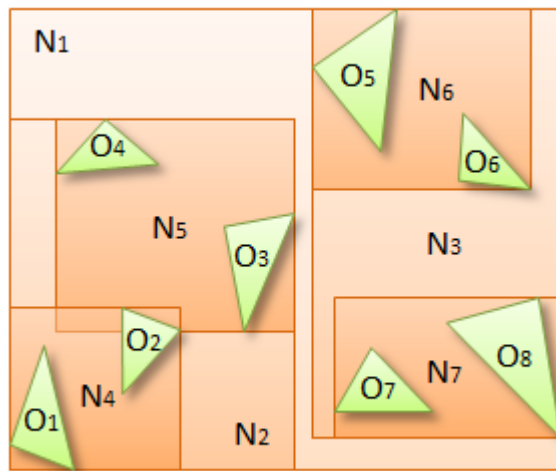


BREAKING 10^9 RAYS/SEC

IAN MALLET

Overview

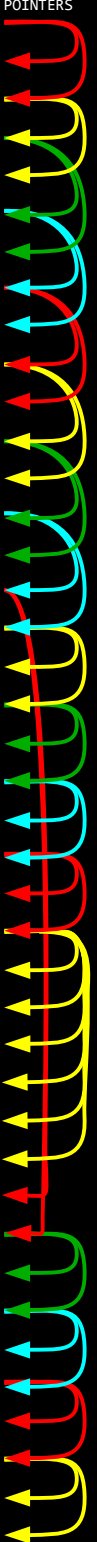


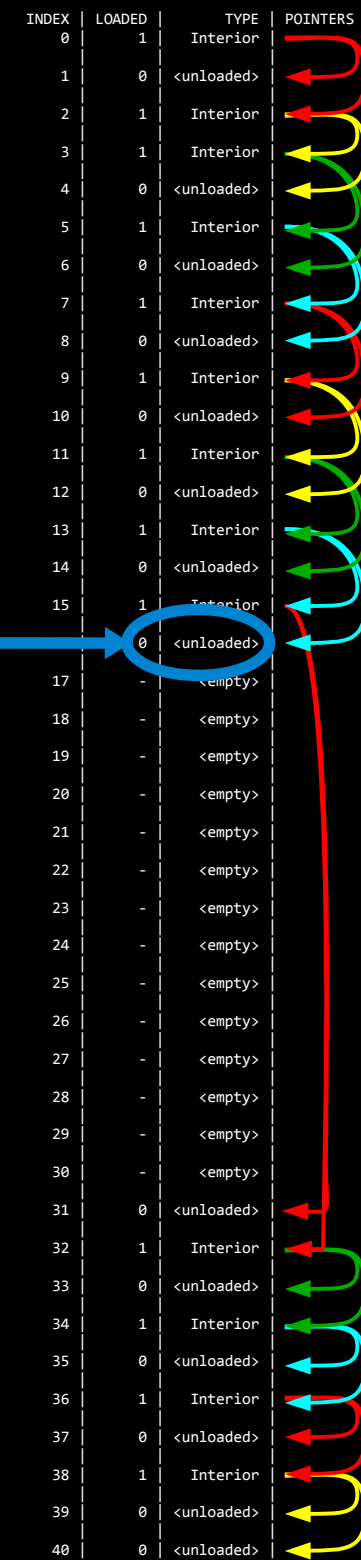
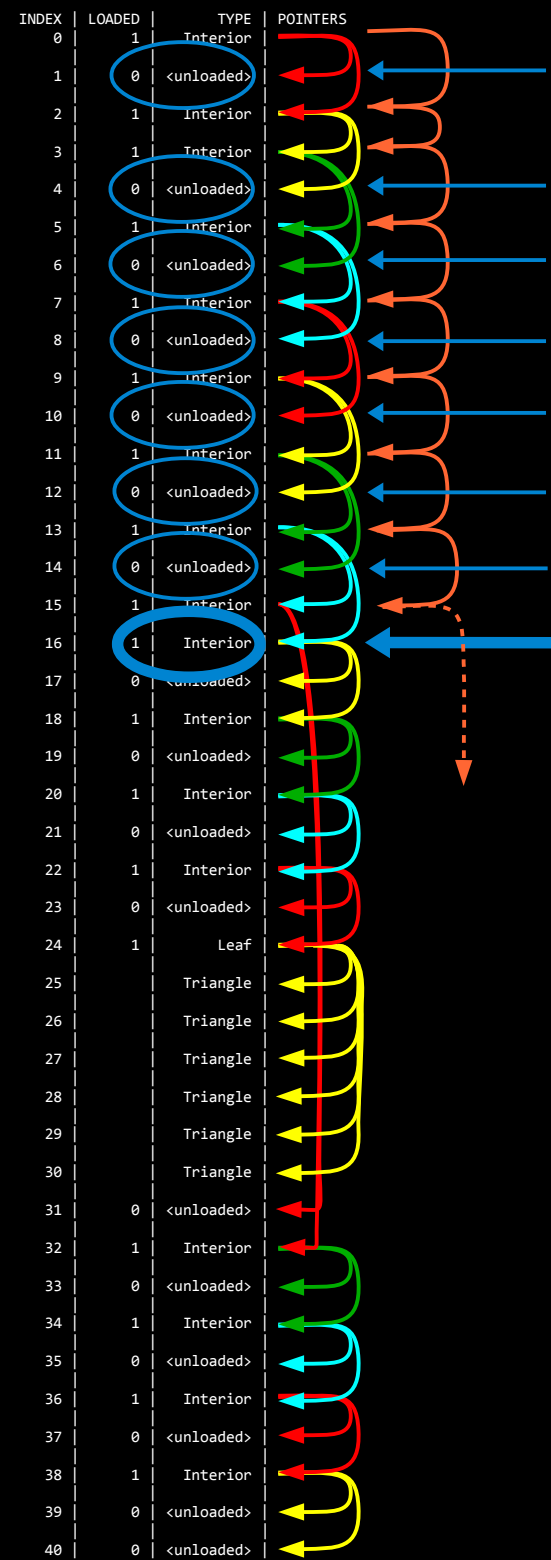
Overview

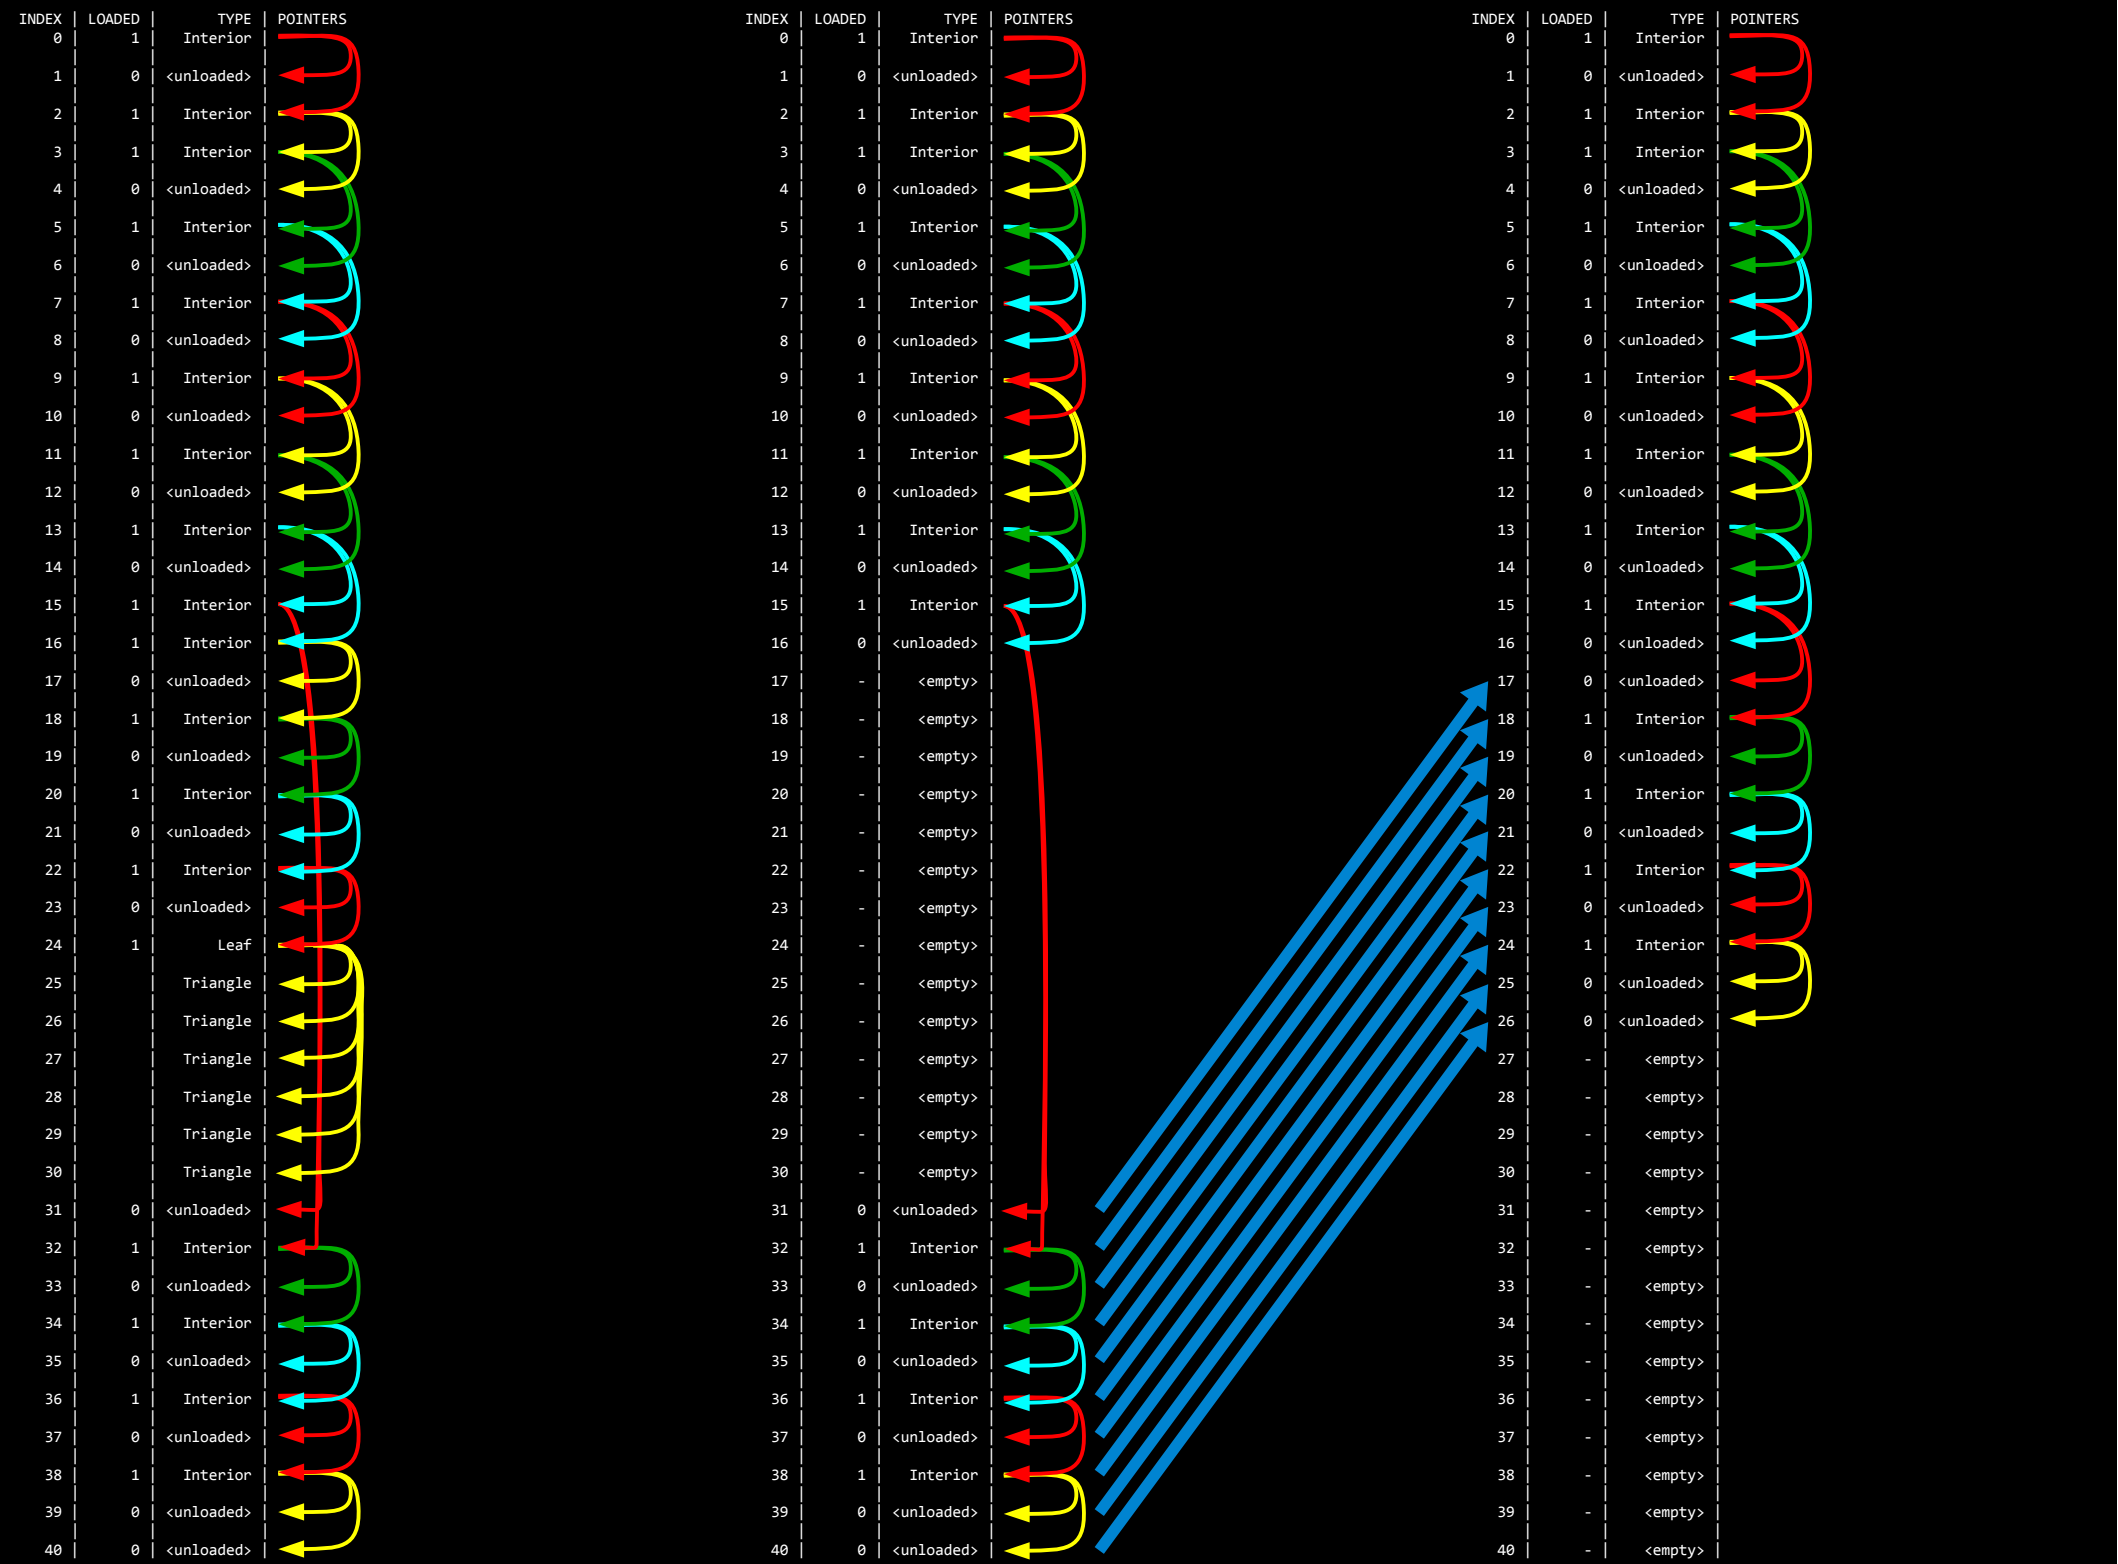
- Goal: gigaray/sec ($1e9$ rays/second) on reasonably sized chip ($< 700\text{mm}^2$, preferably smaller)
- Store current traversal on local stack
- Also as much of several traversals as possible
- **Think of this as a better L1-cache**
- Also optimize stuff/write really good pathtracer

A Full Stack

| INDEX | LOADED | TYPE | POINTERS |
|-------|--------|------------|----------|
| 0 | 1 | Interior | |
| 1 | 0 | <unloaded> | |
| 2 | 1 | Interior | |
| 3 | 1 | Interior | |
| 4 | 0 | <unloaded> | |
| 5 | 1 | Interior | |
| 6 | 0 | <unloaded> | |
| 7 | 1 | Interior | |
| 8 | 0 | <unloaded> | |
| 9 | 1 | Interior | |
| 10 | 0 | <unloaded> | |
| 11 | 1 | Interior | |
| 12 | 0 | <unloaded> | |
| 13 | 1 | Interior | |
| 14 | 0 | <unloaded> | |
| 15 | 1 | Interior | |
| 16 | 1 | Interior | |
| 17 | 0 | <unloaded> | |
| 18 | 1 | Interior | |
| 19 | 0 | <unloaded> | |
| 20 | 1 | Interior | |
| 21 | 0 | <unloaded> | |
| 22 | 1 | Interior | |
| 23 | 0 | <unloaded> | |
| 24 | 1 | Leaf | |
| 25 | | Triangle | |
| 26 | | Triangle | |
| 27 | | Triangle | |
| 28 | | Triangle | |
| 29 | | Triangle | |
| 30 | | Triangle | |
| 31 | 0 | <unloaded> | |
| 32 | 1 | Interior | |
| 33 | 0 | <unloaded> | |
| 34 | 1 | Interior | |
| 35 | 0 | <unloaded> | |
| 36 | 1 | Interior | |
| 37 | 0 | <unloaded> | |
| 38 | 1 | Interior | |
| 39 | 0 | <unloaded> | |
| 40 | 0 | <unloaded> | |



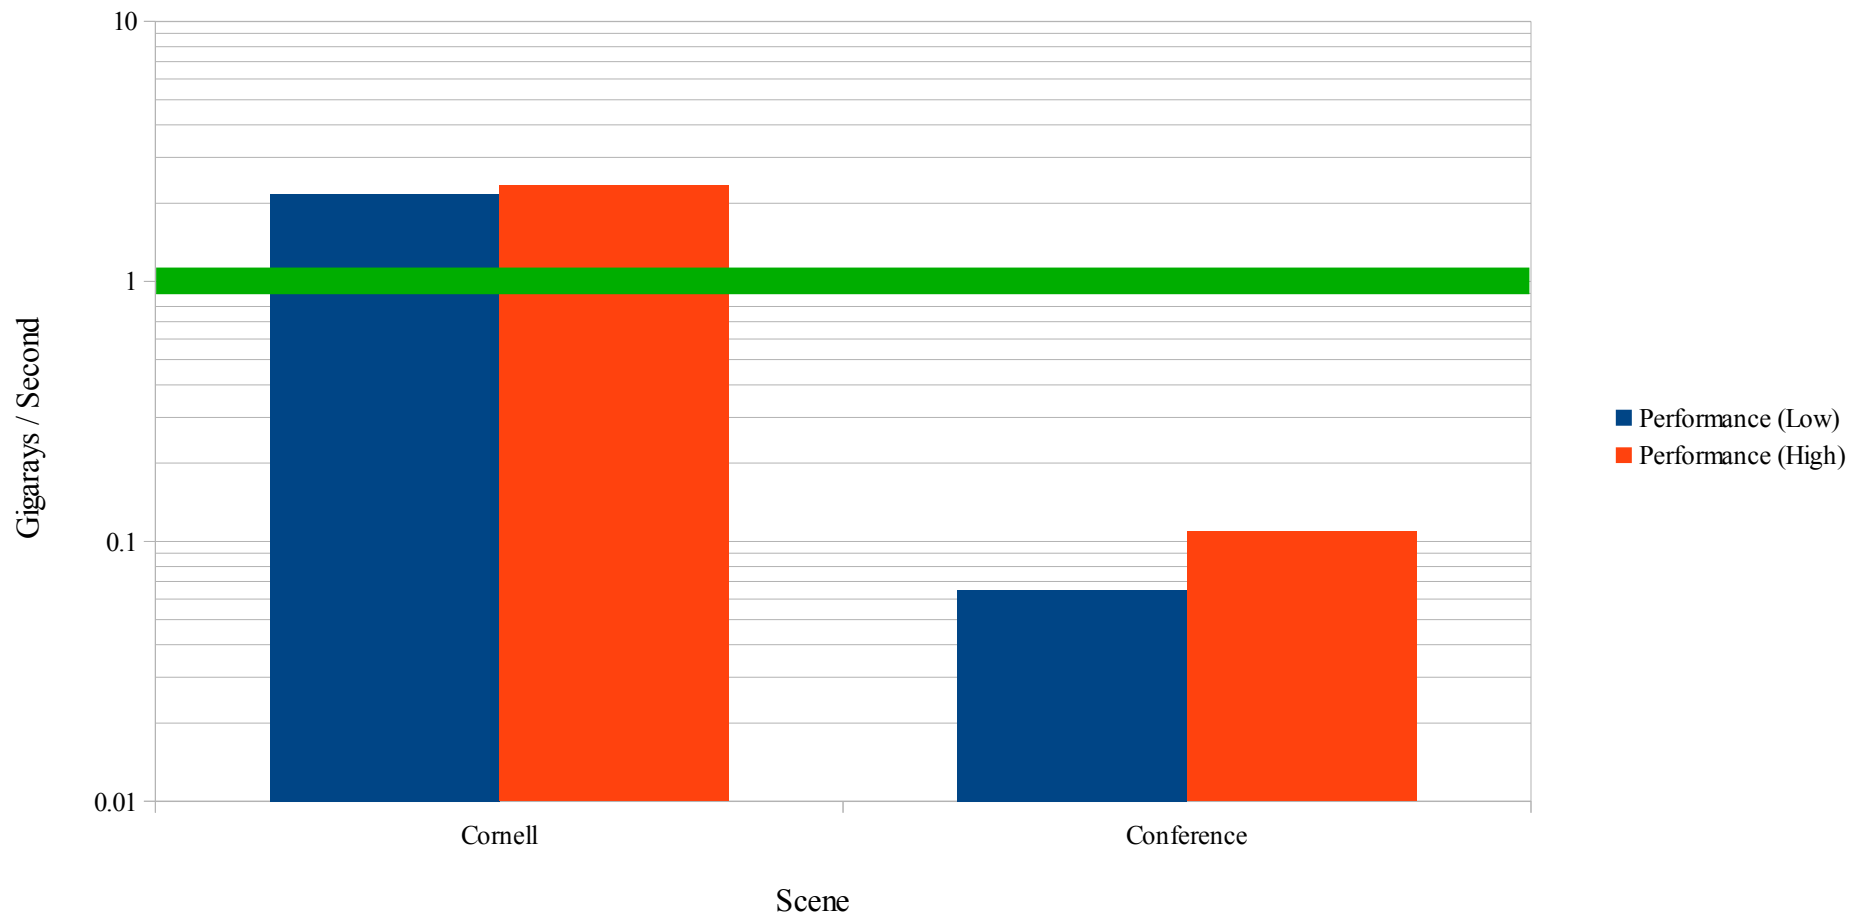




Results (Major Factors)

- A lot of scene fits on the stack (reuse)
- Shadow rays (coherent, occluders usually nearby)
- Very little of scene fits (less reuse)
- Threads jumping in screen-space with shallow depth (incoherency, less reuse)
- Motivating example . . .

Results (Min and Max)



Results

- Cornell
 - Everything fits on stack, and each processor can reuse
- Conference
 - Not everything fits on stack and processors jump in image space a lot → much less reuse
- Hairball/Dragon (simulation too large → too long)

Performance vs. Area

- Area calculation bogus for localstore size (constant)
- Used default size (with allowances for simulator/debug info.), map to same size stack
- $297\text{mm}^2 \rightarrow$ reasonably sized chip!
 - Inaccurate; but probably not by too much

Performance Comparison

- Reference path tracer runs . . . differently
 - Use L1 and no L2
 - Try to match areas
 - Unsuccessful; and area calculations bogus anyway
- Certainly looks like my way is faster overall
- Weird slowdowns though.
Might be ray profiling code.

Difficulties

- Simulator
 - Very slow
 - No debugging
 - Bugs and unimplemented features
- Clever mistakes
 - ~2,000 lines of deep pointer arithmetic
 - Fancy compiler options (gah)

Conclusion

- Mixed results
- Should be combined with ray rescheduling of some kind (even just one thread gets 16 pixels → big win)
- Lots more comparison work is needed → weakest part of project
- *Still* room for optimization . . .

Questions

Image Credits

- me
- others