

# Practical Parallelism

Tales from the trenches of commercial place and route tools

Adrian Ludwin, Altera Corp.

ParCAD @ ICCAD 2010

San Jose, November 11, 2010

*Additional information in PowerPoint speaker's notes*

# Purpose of today's talk


- Warn you about problems we encountered
- Encourage researchers to study certain topics
- Suggest ideas for tool vendors

# About myself

- Software engineer on Quartus II
  - Altera’s CAD suite for FPGA design
- Worked on several parallel algorithms in QII
  - Either developed them myself or assisted other engineers
- The hardest algorithm was parallel placement
  - A Ludwin, V Betz and K Padalia: “High-quality, deterministic parallel placement for FPGAs on commodity hardware,” FPGA 2008

# Why Altera is investing in parallelism

*Here's that slide again!*



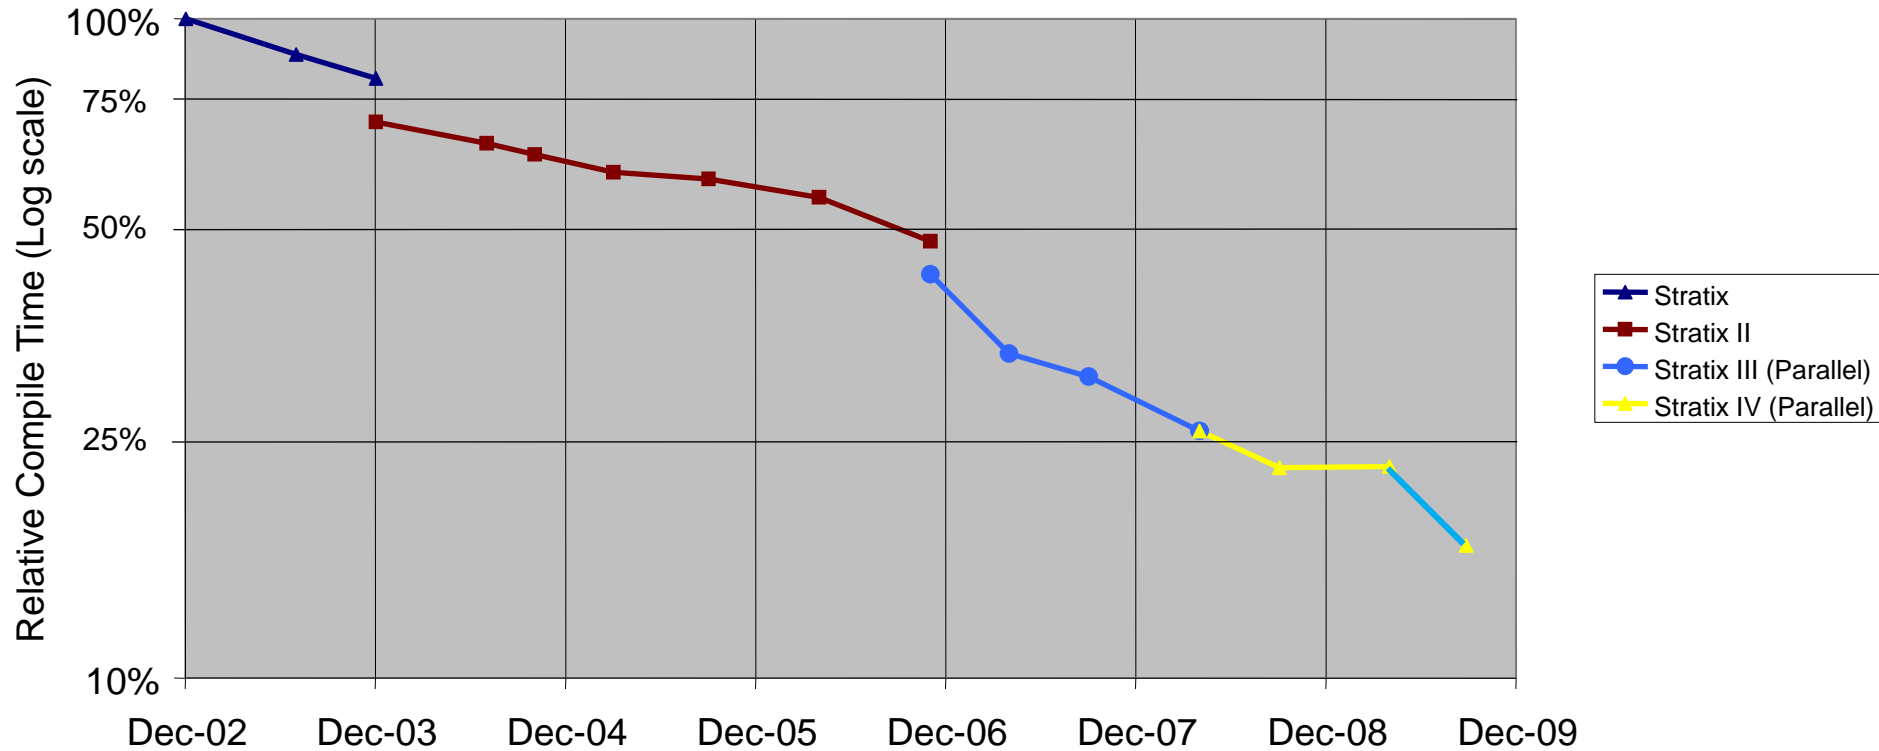
© 2010 Altera Corporation—Public

ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS & STRATIX are Reg. U.S. Pat. & Tm. Off. and Altera marks in and outside the U.S.



# Quartus II compile time history

## *Fixed design on a fixed CPU*



***Not shown: savings from hierarchical/incremental design***

# Agenda: three types of challenges

- Development
- Determinism
- Quality

# Development challenges

# Who's embarrassed?


- QII has few “embarrassingly parallel” algorithms
- Many algorithms have thousands of lines of code
- Fitter (“P+R”) has many distinct algorithms
  - Periphery placement (eg PLLs, transceivers, clocks)
  - Register packing
  - Assorted netlist transformations
  - Clustering
  - Placement
  - Congestion analysis
  - Slack allocation
  - Routing
  - Timing analysis

© 2010 Altera Corporation. High-speed/low-power tile setting

ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS & STRATIX are Reg. U.S. Pat. & Tm. Off.

and Altera marks are trademarks of Altera Corporation in the U.S. and other countries.

# Fitter: no one major runtime hog!



*One Stratix V design, internal 10.1 build. Total Fitter time: about three hours*

© 2010 Altera Corporation—Public

ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS & STRATIX are Reg. U.S. Pat. & Tm. Off. and Altera marks in and outside the U.S.



# Development challenges

Part 1 of 3: ensuring thread safety

# About Quartus II code

- Mainly C/C++
- Most algorithms have similar template:
  - Initialize permanent data structures
  - Initialize temporary (“scratch space”) data structures
  - Do work
  - Remove data structures

# That's easy!

- Simple to make thread-safe, in theory
- Just find all global variables (in C):
  - Permanent data: lock or make per-thread copies
  - Scratch space data: make per-thread
  - Only care about data that's *written*, not read-only
- Harder in C++ but similar idea
- Problem: hard to find suitable tools

# Tools to find possible data races

- Existing tools are overkill for our simple problem
- Intel Parallel Advisor?
  - Didn't exist until recently
  - 40x runtime
  - 12x memory
  - Requires testcases
  - Emits message for every *error* instead of every *variable*
- Coverity?
  - Harder to run
  - Expensive (if you don't already have it)
  - Also emits messages for errors instead of variables

# So we wrote our own

- GCC backend, Perl/Tk front-end
  - GCC finds functions, globals
  - Post-processor reads GCC output
  - Also maintain list of “known good” functions/globals
- We called it “Deep Code Inspection”
- Anyone want to develop it for us?

# Development challenges

Part 2 of 3: writing the algorithm

*Covered in the rest of this presentation*

# Development challenges

Part 3 of 3: quality assurance

# Testing issues

- Runtime
- Stability
- Determinism

# Testing runtime

## *Options on a twelve-core box*

### Run 1 instance of QII

- Very poor throughput
- Near-perfect runtime measurements

### Run 12 instances of QII

- Good throughput
- Noisy, low-fidelity runtime measurements

# Testing stability

- Now, we use standard tools
  - Coverity, Thread Checker, helgrind
  - Push runtime tools to/over to the limit
- We allow “unsafe” operations by design
  - Our algorithms are allowed to detect these and recover
  - But the tools have no way of knowing this
  - No good way to separate “false” positives from real mistakes
- We rely heavily on our own stability testing

# Testing determinism

- Compare final output of multiple runs
- Use “cookies” to further ensure determinism
  - Can’t just look at the final result
  - What if we chose *not* to do something?
  - Did we make that decision deterministically?
- Difficult to choose good testcases
  - Running them all multiple times would take thousands of hours

# Development challenges summary

- Thread safety is hard, but tools could help
- Testing is hard, and we just have to live with it

# Ensuring determinism

# Is this really important?

- Rarely studied in academic literature
  - Eg: of the ~10 papers I cited, only two were deterministic, and one of them didn't even bother to mention it
- But we think it's essential
  - Both for the vendor and the customer

# Motivation for vendors: quality

- Deterministic tools are easier to verify
  - Especially given our difficulty with tools
  - Testing for determinism achieves similar goal
- Reproducing bugs from customers

# Motivations for customers

- Customers do re-run unmodified designs
  - Not good if the result changes randomly
  - Altera sometimes has to re-run their designs as well
- Some customers *require* determinism (military)
  - Concerned about software being compromised
  - Need to compare output to a “known good” version
  - If output changes randomly, this is not possible

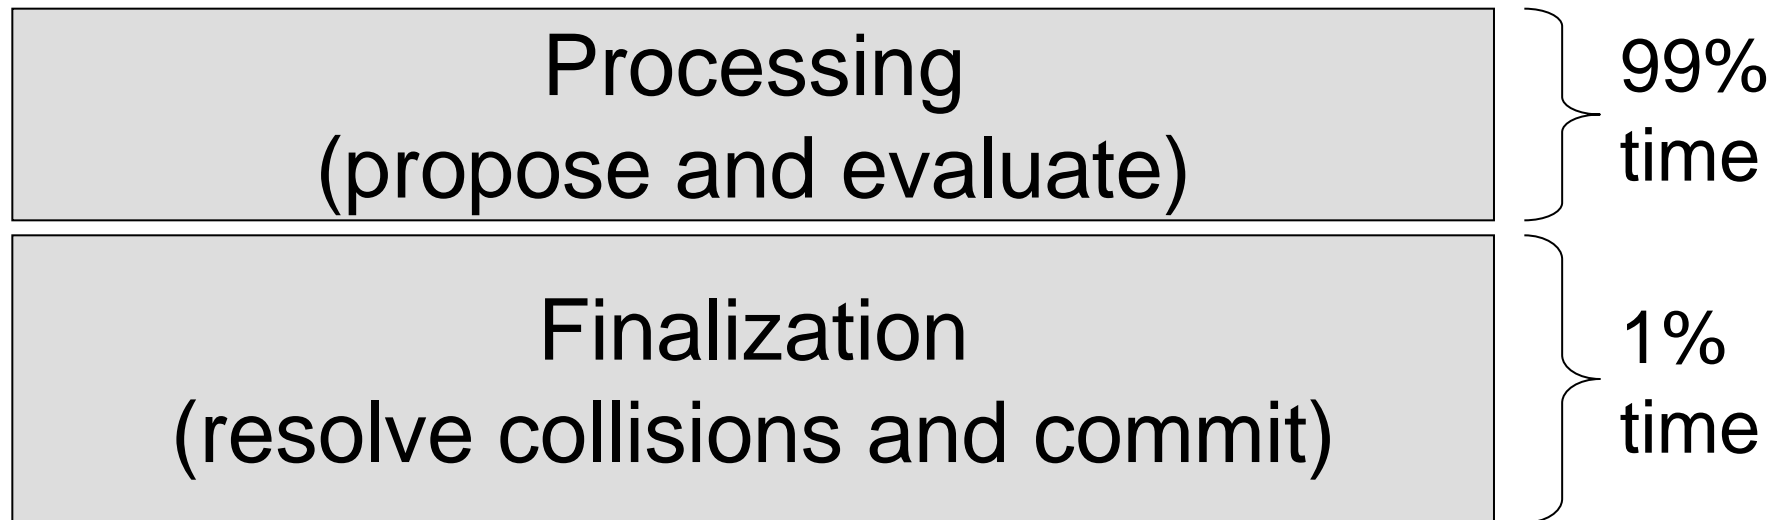
# Challenges achieving determinism

- Algorithmic development
- Algorithmic performance
- Hardware limitations

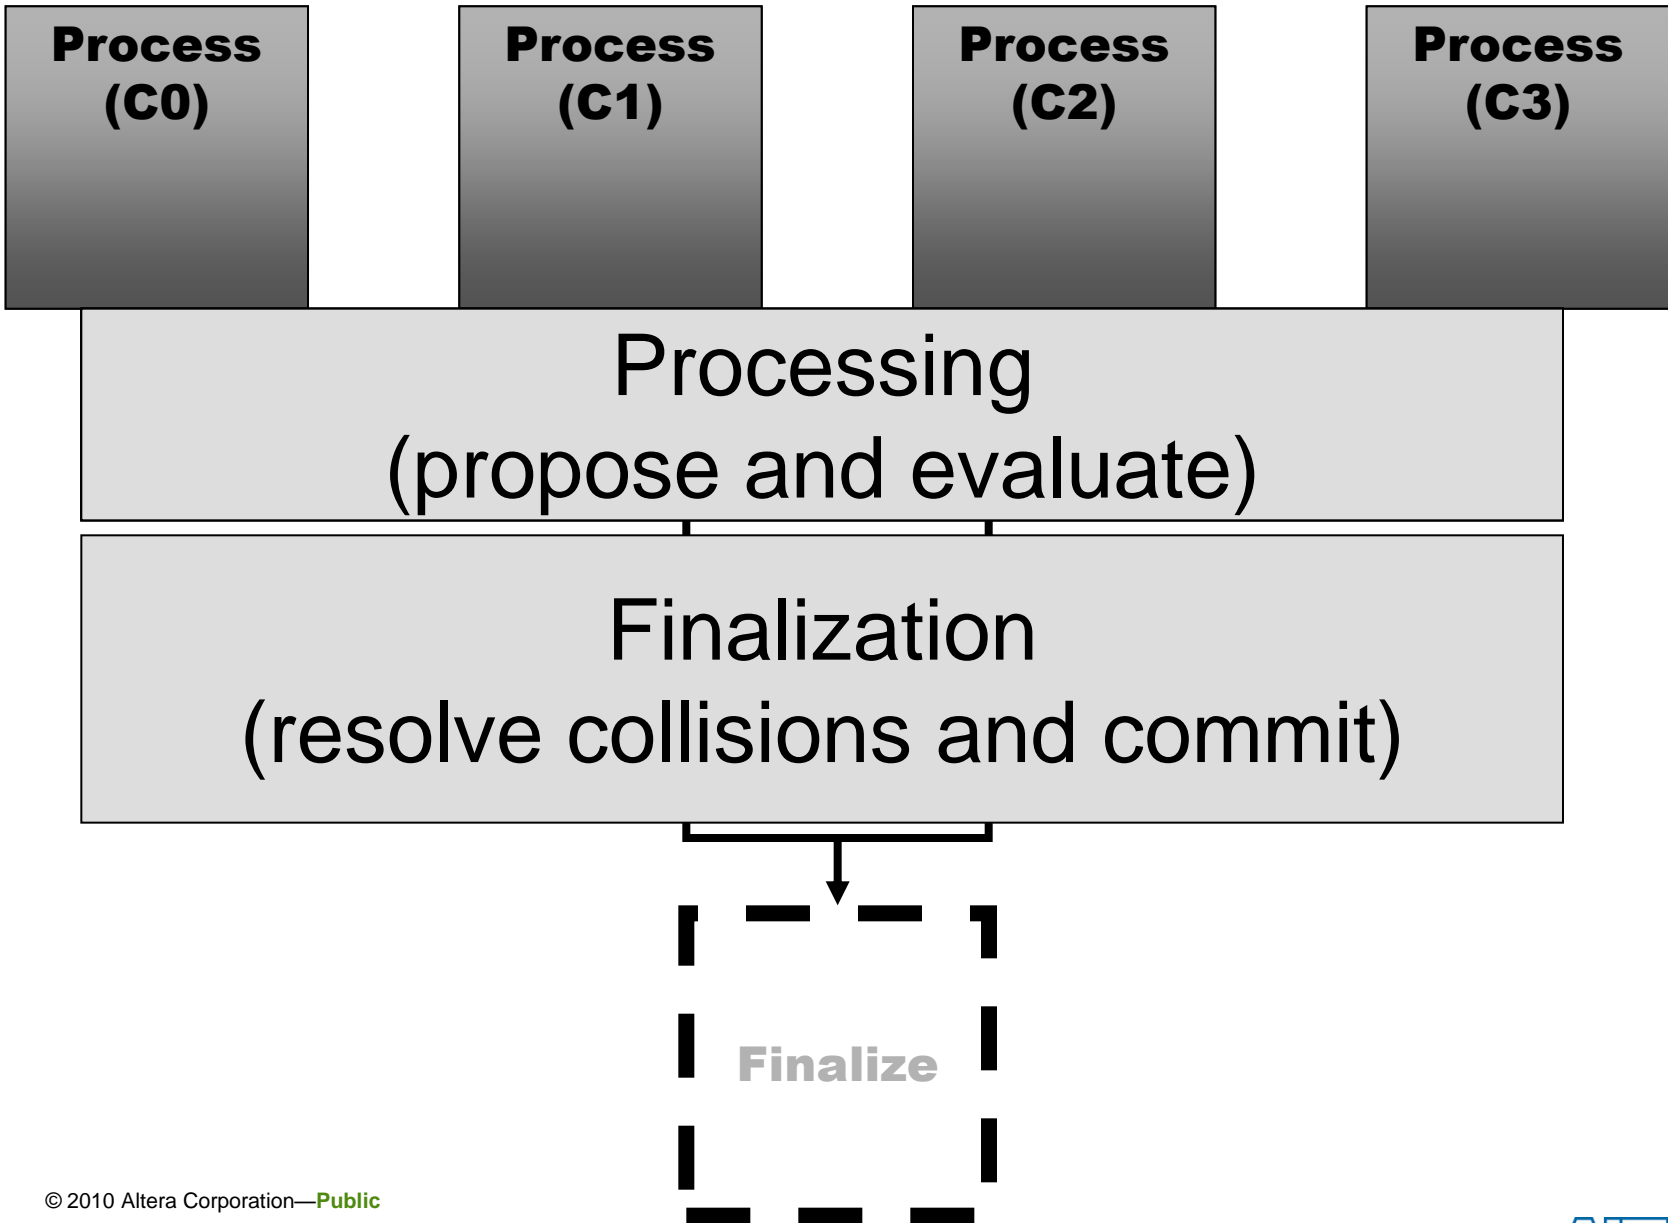
# Challenge: algorithmic development


- First approach: strictly partition the problem
  - *Guarantee* no interactions between cores
- Second approach: track and resolve interactions
  - Compatible with a partial partitioning
  - Are you *sure* you've actually tracked/resolved everything?
  - Building an algorithm to resolve interactions can be difficult

# Building a deterministic placer



This is the inner loop of the Quartus II placement algorithm. We divide it into two stages: “processing” and “finalization.” The “processing” stage is the one that is parallelized.






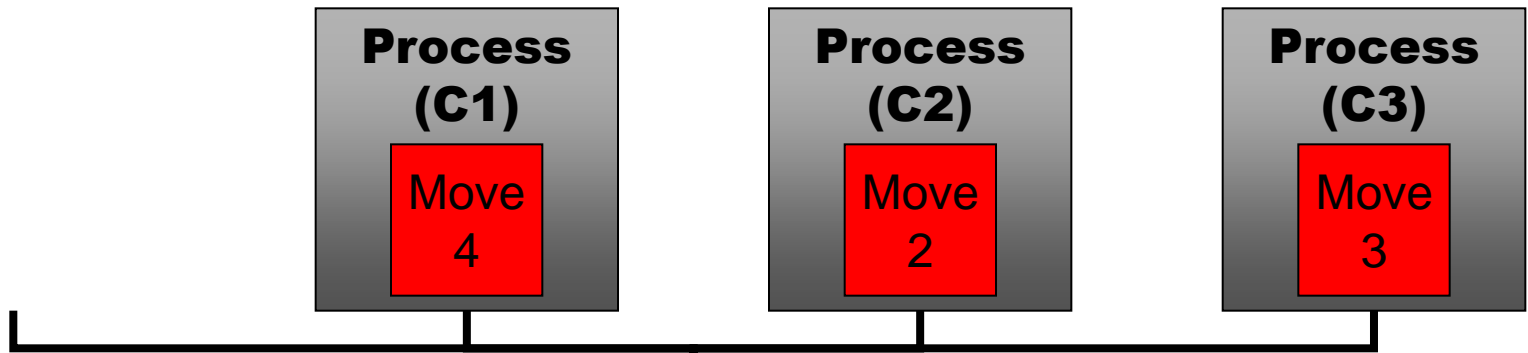
Core that processed this last move becomes responsible for finalizing moves in the queue. Note that it does not have to wait for any other core; it simply inserts the move it inserted to the front of the queue.

When a core finishes out of order, it sits in the queue until the earlier moves are finished. Meanwhile, the core that processed it goes onto the next move. Cores do not stall and wait for any other cores.




All four cores begin processing moves at the same time. Since finalizing moves is so fast, it would be a waste to devote a core to that task. Instead, all cores have the ability to finalize moves at the appropriate time, as this example will show.

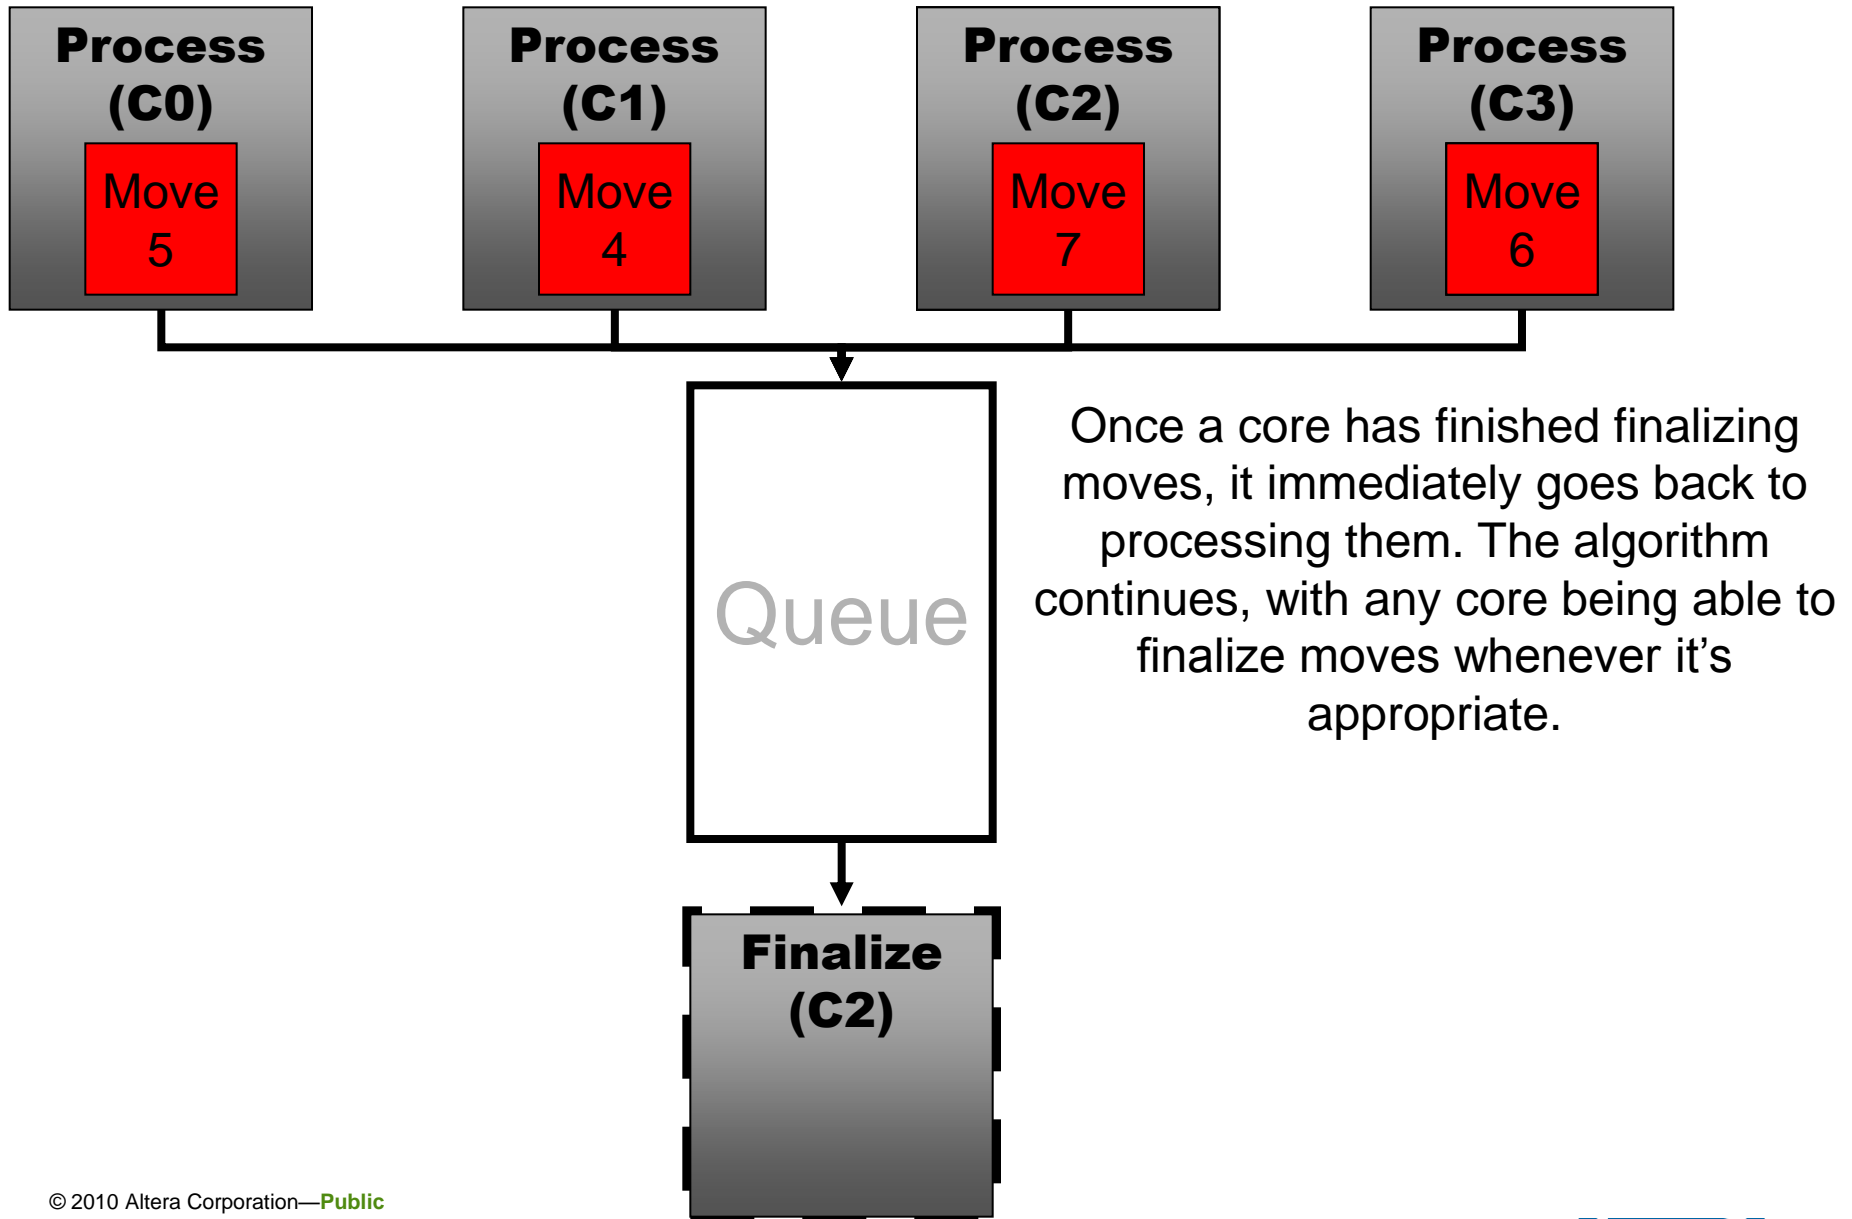
The priority queue now has two moves ready to be finalized.



The core that processed this last move now becomes responsible for finalizing all the moves in the queue. Note that it did not have to wait for any other core; it knows that the move it inserted went to the front of the queue.



Finalization is where all interactions between moves (aka collisions) are resolved.



# Challenge: algorithm performance

- Partitioning often costs runtime or quality
  - Quality: cross-partition modifications forbidden or discouraged
  - Runtime: maybe every partition didn't need same optimization
- Tracking/resolving costs runtime
  - Both tracking and resolving incur direct overhead

# Challenge: hardware limitations

- Determinism implies intra-core communication
- Intra-core communication is slow
  - Limited bandwidth between cores
  - Commonly-used data is not in your cache
- More fundamentally: hard to distribute

# Possible solutions


- No silver bullet, of course
- Hardware transactional memory: bronze bullet?
  - Pros: moves much of the overhead to hardware
  - Cons: doesn't actually exist for the foreseeable future

# Determinism summary

- We think it's vital
- It's hard to develop
- It has performance implications

# Parallel Quality

# We all know the runtime-quality tradeoff...




© 2010 Altera Corporation—Public

ALTERA, ARRIA, CYCLONE, HARDCOPY, MAX, MEGACORE, NIOS, QUARTUS & STRATIX are Reg. U.S. Pat. & Tm. Off. and Altera marks in and outside the U.S.




# ... and this is what we want:




# Unfortunately, it's not that simple

- Parallel likes lots of cleanly divided subtasks
  - Traditional “fast” algorithms didn't really care
- But many CAD algorithms hate partitioning
  - Can reduce both quality and runtime
- Other simplifications can have the same effect


# So instead of this...



... we “simplify” the algorithm to this ...



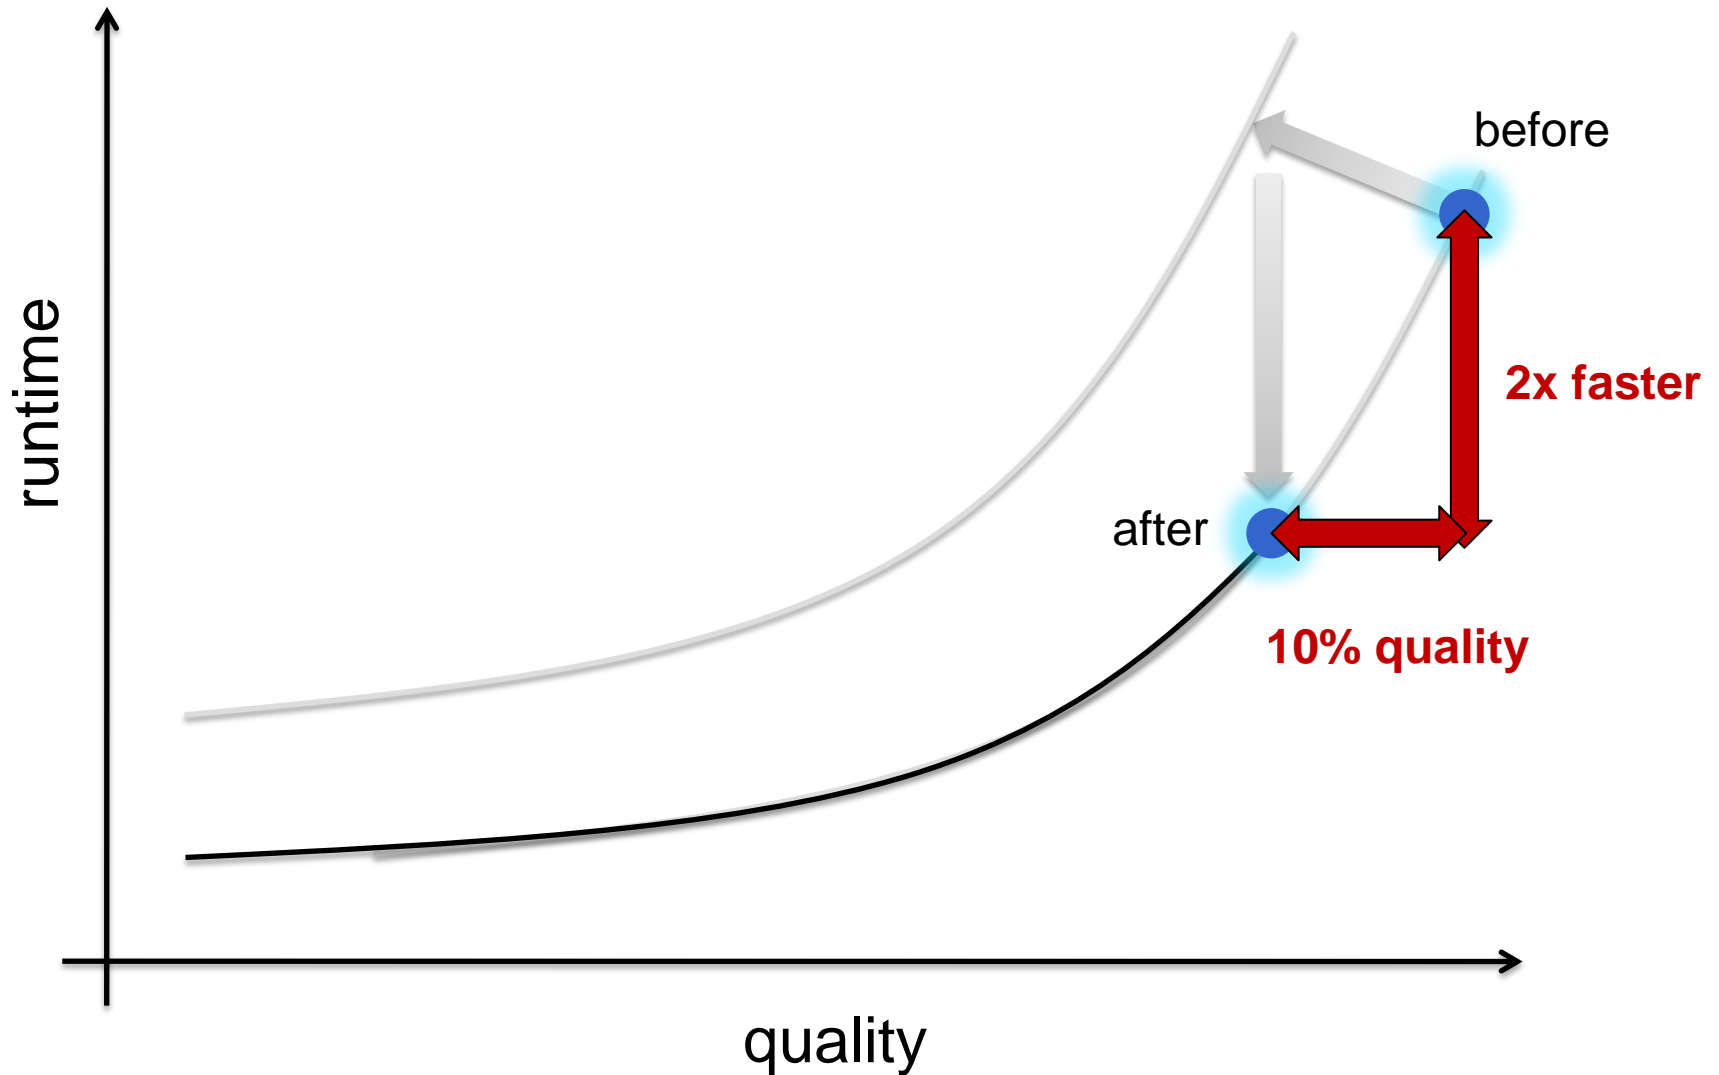
# ... well that was fun, wasn't it?



# This is sometimes easy to miss

- Parallel algorithm: “2x faster, for 10% more wire”
  - Please note that the example that follows is an exaggeration


# This is sometimes easy to miss




# What are the alternative algorithms?

- What if you only consider the simpler algorithm?
  - You might achieve a complex algorithm's runtime/quality curve
  - So it's important to know the alternative algorithms
- Could be better to accept a lower speedup
  - Obviously this depends on how much speedup is on the table

# Lower speedups are sometimes better



# Lower speedups are sometimes better



# Concrete example: simulated annealing

- Inner loop: propose move and evaluate cost
  - A cell and proposed new location is selected at random
  - Proposed location is generally close to present location
  - Costs are generally local: eg, nets connected to the cell
- This is *relatively* easy to parallelize
  - Uniform spatial distribution of work with good spatial locality
  - Easy to partition
  - Eg Sun + Sechen '97: near-linear speedup on 6 machine cluster
    - No increase in wirelength
- Example: VPR placement algorithm
  - An academic FPGA place-and-route package

# Complication: Quartus II placement

- Based on VPR placement algorithm
- Many important improvements over VPR
  - Neither cells nor locations selected at random (“focused moves”)
  - Proposed locations may be distant from current location
  - Many costs are less local: eg, path-based timing, clock regions
- Less spatial locality - harder to parallelize
  - 2.4x average speedup on 8 cores during the anneal (up to 3.7x)
  - 4.0x average speedup on 8 cores during the quench (up to 5.6x)


# Comparison: VPR vs Quartus II

- At default optimization levels<sup>1</sup>:
  - 15% faster routed critical path, 7% lower routed wireuse
  - 35% lower placement time
- VPR can't match QII's quality for any runtime
  - Eg increased VPR runtime by 10x: still 12% worse critical path
- A parallelization of VPR that doesn't work on QII is of little use to customers who need advanced timing closure

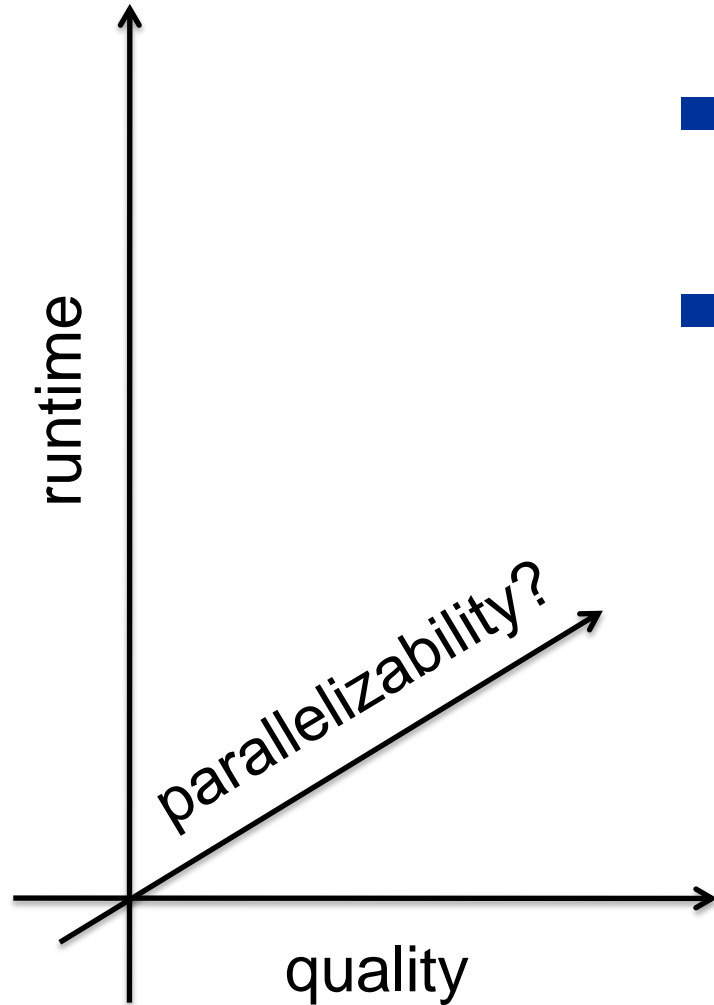
*If you're wondering why we use simulated annealing at all, ask me after the talk...*

<sup>1</sup>Based on internal testing. Quartus II is very difficult to compare fairly to academic tools.

# So are we stuck with this?



# Parallelizability: a third dimension?



- Not entirely independent of quality or runtime
- But there may be optimizations that don't have a large impact on quality or (serial) runtime

# Example: Quartus II router

- Dynamically assigns nets to processors
  - A type of partitioning
- No loss of runtime or quality when run in serial

# Example: Quartus II placement?

- Currently: we do not partition
  - Therefore, every move can interact with every other move
- Future work: *biased* moves
  - Cores would *tend* to work on cells in physical regions of the chip
  - Therefore moves *tend* not to interact
  - Can simplify interaction tracking/resolving in many cases
- Hopefully will improve parallelizability without otherwise impacting runtime or quality

# Quality summary

- Be careful what algorithms you parallelize
- Be careful of simplifications that impact quality
- High-quality algorithms may be harder to parallelize for less speedup
  - But overall quality might still be better

# Conclusions

# Future work

- We continue to face three major challenges:
  - Development issues (eg tools, testing)
  - Ensuring determinism
  - Balancing parallel speedups with quality
- Good news: fascinating research opportunities
  - Also: parallel CAD software engineers remain gainfully employed
- Are you a parallel researcher?
  - You may find these topics interesting and useful to investigate
- Are you a tool vendor?
  - Maybe this gives you some ideas for future products

# Thanks!

Adrian Ludwin - [aludwin@altera.com](mailto:aludwin@altera.com)