FMCAD Panel:
Model Checking in the Cloud

Olivier Coudert
SiCAD Inc.

October 25th, 2012

SiCAD™
CLOUD-AIDED SILICON DESIGN

Topics

- Cloud computing
- Distributed model checking
- Challenges

Cloud Computing Promises

- On-demand computing resources
- No upfront costs
  - pay as you go
- Scalable
  - 100's of cores assembled in a compute grid
  - TB's of storage
  - 1Gbps LAN, 10Gbps HPC
- Expand geographic reach

SiCAD™
CLOUD-AIDED SILICON DESIGN

Performance Scaling

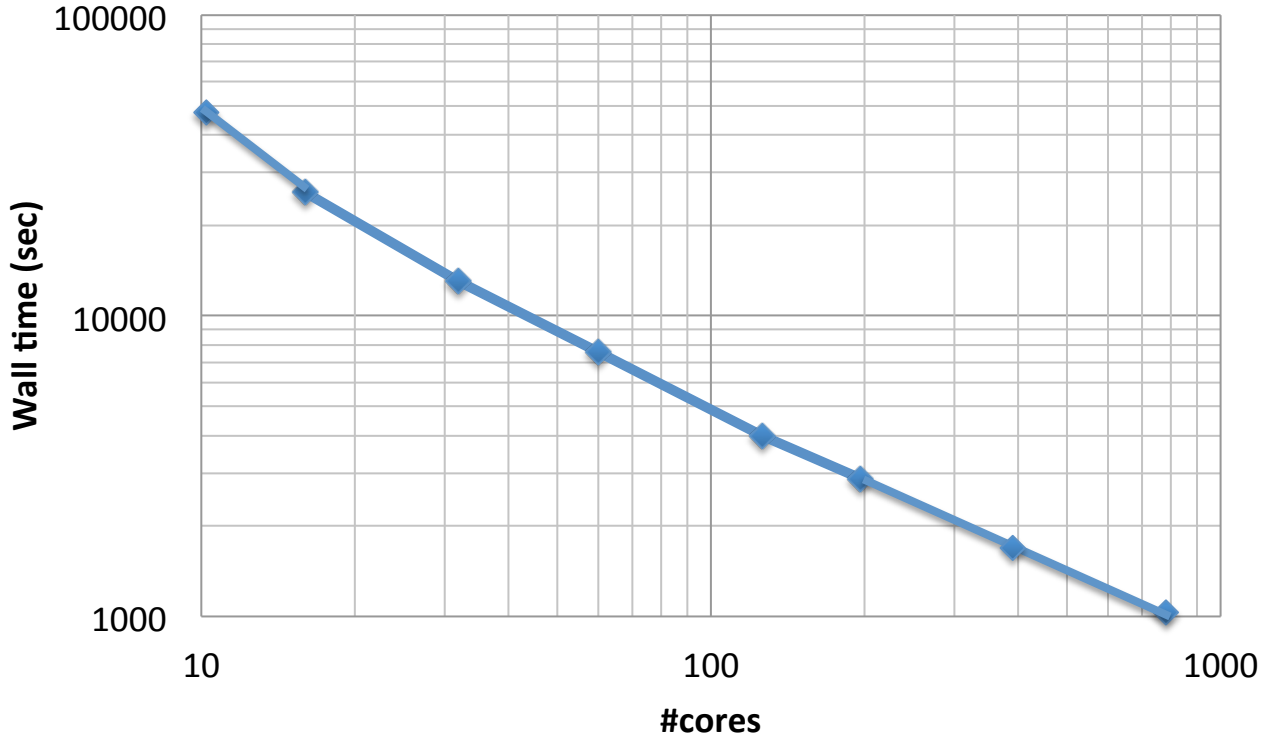

- Cluster setup time : 10-15mn
- Application: physical verification
- 10 cores: 13h42mn
- 768 cores: 17mn

SiCAD™
CLOUD-AIDED SILICON DESIGN

Distributed Model Checking

- Parallelism has many flavors
- In practice: MIMD
  - Network of machines
  - Distributed memory with multiple cores

- Model checking
  - LTL, CTL, etc
  - State exploration

Explicit State Exploration

- Explore state one by one
  - DFS or BFS state exploration
  - Need to recognize visited states
  - Mostly memory limited
- Parallelization
  - Partition state space, and assign each partition to a node of the grid
  - Partition: hashing, windowing

Implicit State exploration

- BDD-based
  - BFS state exploration
  - Mostly memory limited
- Parallelization
  - Partition variables, and assign each partition to a node of the grid
  - Partition made of consecutive variables
  - BDD node management is breadth-first
  - Distributed hash-tables for BDD operations caches

Bounded Model Checking

- SAT-based
  - Unroll model k times
  - Mostly time limited
- Parallelization
  - Partition Boolean space (assume some variables have some constants values)
  - Conflict clauses need to be shared

Cloud Models

- Private cloud managed by EDA vendor
  - Aldec (logic simulation)
  - Nimbic (3D simulation)
  - Tabula (FPGA synthesis)
  - Cadence (reference flow)

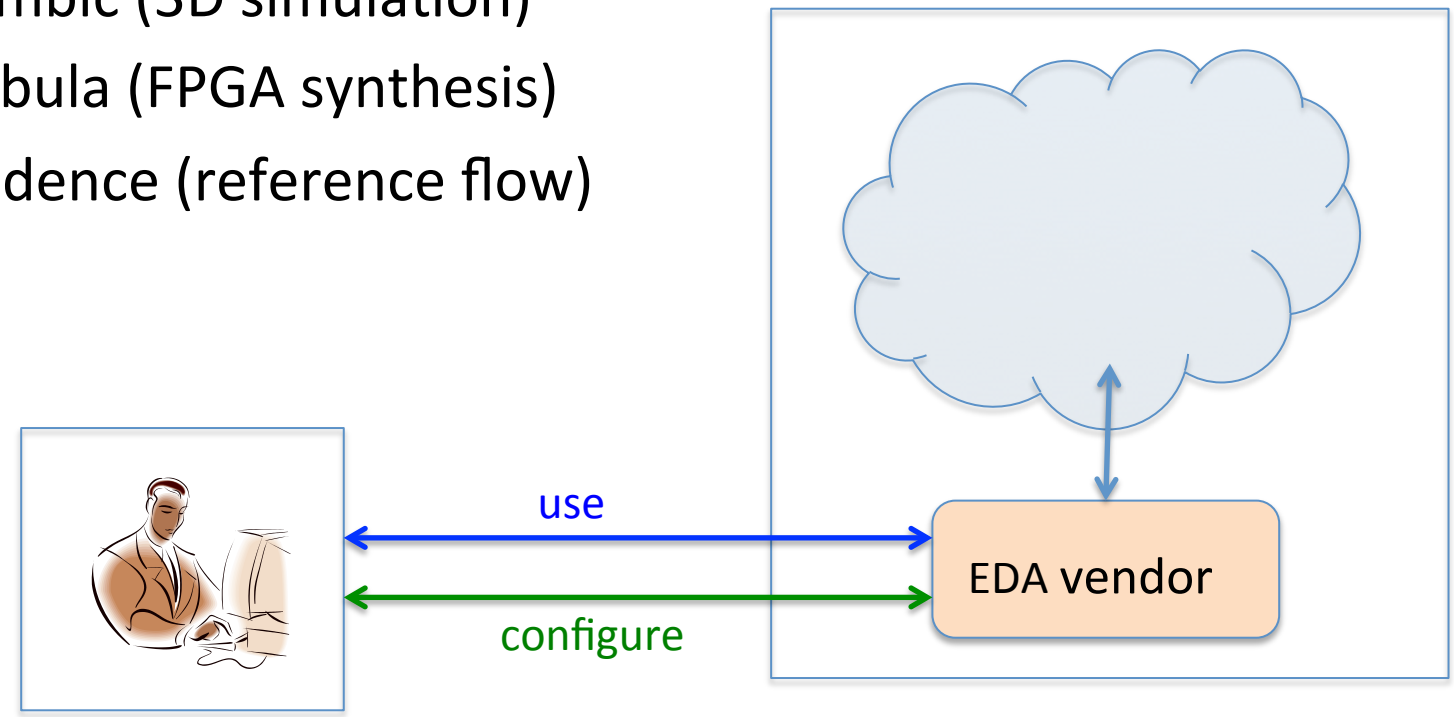

Cloud Models

- Public cloud configured by EDA vendor
  - Synopsys (logic simulation in AWS)

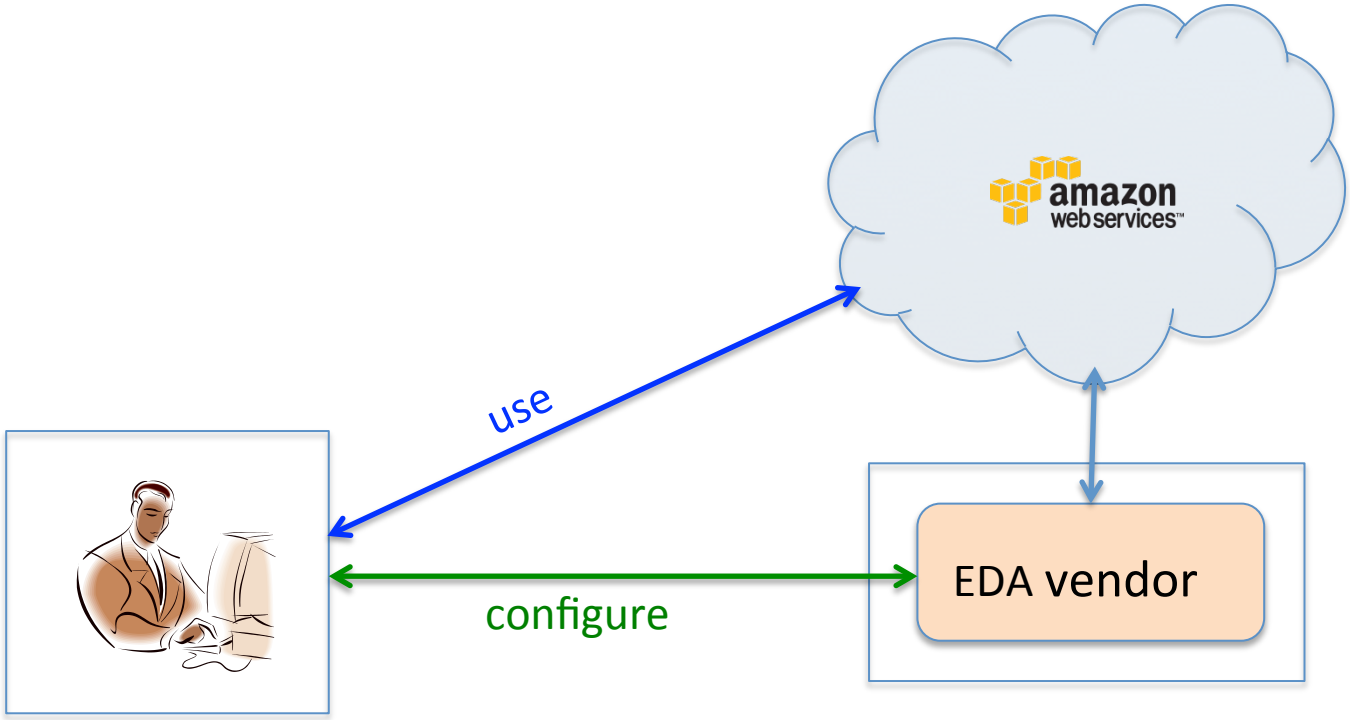

SiCAD™
CLOUD-AIDED SILICON DESIGN

Cloud Models

- Cloud platform configured and managed by a 3rd party
  - Xuropa (SW evaluation in AWS, used by Synopsys, Cadence, and Xilinx)
  - Plunify (FPGA synthesis in AWS)
  - SiCAD

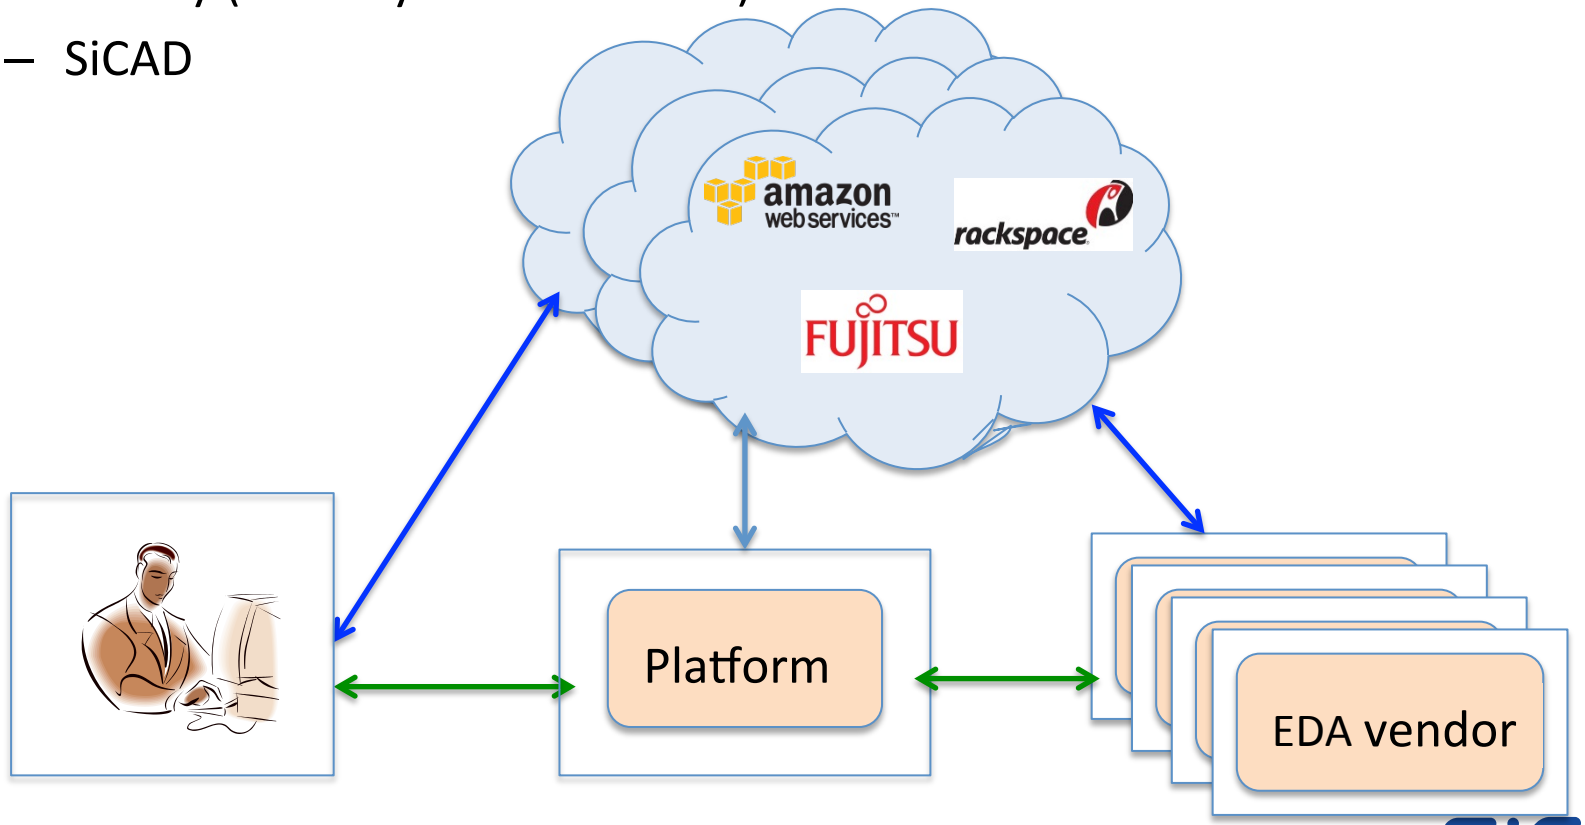

SiCAD™
CLOUD-AIDED SILICON DESIGN

Challenges

- Legal
  - SLA
  - Liability in case of data loss or breach
  - Geographical location of data
  - Cloud provider origin
- Multi-party agreement
  - Multiple EDA vendors, design house, foundry, cloud provider
- Business model
  - SW needs a pay-as-you-go model
  - Risk to cannibalize TBL's revenue for EDA vendors

Challenges

- Technical
  - Scalability of application
  - Fast, fault-tolerant, compute grid provisioning and setup
  - Volume of data transfer
    - 10GB @ 30Mbps: 44mn
    - 10GB @ 1Gbps: 1mn20sec
- Security
  - Highly sensitive data (design, SW, and IP)
    - Data confidentiality –transmission, at rest
    - Data integrity –e.g., disaster recovery
    - Data availability –uptime, latency
    - Data disposal –data removal and storage disposal
  - Customer may want to keep its SW usage confidential

Rethink for distributed in the cloud

	1Gpbs LAN	Hard drive	SSD	RAM
latency	0.5ms datacenter roundtrip	3-10ms	0.1ms	100 ns
bandwidth	128 MB/s	140 MB/s	100-600 MB/s	6-17 GB/s
capacity	N/A	up to 8TB	256GB - 1TB	4-64GB
cost	free	$0.05/GB	$0.65/GB	$5-10/GB

- Writes are expensive, reads are cheap
  - Once read, data is cached
  - Writes are ~50x slower than read
- It might be faster to move data chunks in the LAN than reading it from a hard drive
- SSD is changing the way data can be managed

SiCAD™
CLOUD-AIDED SILICON DESIGN

Conclusion

- Cloud computing
  - Large, cheap, readily available compute grid
- Model checking
  - Need algorithms that can leverage a large distributed computing network (100-1000+ cores)
  - Licensing needs to follow burst computing models
  - Security is a bottleneck

SiCAD™
CLOUD-AIDED SILICON DESIGN