




Limits to Growth?

Gerard J. Holzmann




evolution of the desktop PC



time to fill 10 GB of RAM

10⁶ states per second; 10⁹ states per second; 10¹² states per second
(up to 10¹² states in bitstate mode; upto 10⁹ states in normal mode)



2/2/08

3

scaling

- in **one day** we can explore no more than 10¹⁰ states
 - no matter how much RAM or disk we have!
 - if Moore's curve for clock-speed had continued, we would have continued to expand this range
- using **multi-core**, in the best case we could increase our range by using multiple CPU cores
 - but concurrency != scalability
- some verification problems are **larger** than what we can handle (and not amenable to symbolic methods or abstraction)
 - how do we handle those?
 - the infinite state space and the infinite memory




2/2/08

4

spin version 5



- supports multi-core verification
- developed on a dual quad-core system with 32 GB of memory
- linear scaling is achieved in the best cases measured




2/2/08

5

multi-core model checking

leader election (N=9), reference model, and a phone switch model




0.5 M states

2/2/08

6

beyond 8 cores: and apparent anomaly




what happened to the nice linear scaling?
hypothesis: are *memory caching protocols* getting in our way?

2/2/08

7

memory access on SGI Altix with fast NUMA interconnect

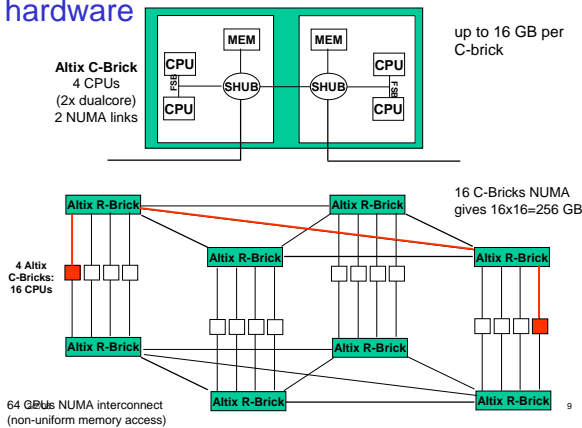


Source: Scalability: The Software Problem
Jonathan Appavoo, Volkmar Uhlig, Dilma da Silva,
Proc. STMS'07, San Jose, CA, March 2007.
Second Workshop on Software Tools for Multi-Core Systems

2/2/08

8


hardware



9

a simple experiment

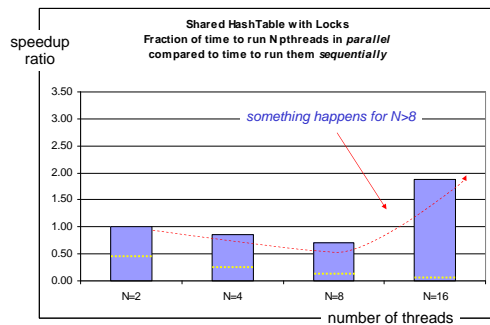
- a small test program that writes S "states" of V bytes each into memory
 - the program simulates the actions of a model checker: randomly generating states, computing hashes, and storing the state in memory
- execute this program as N parallel threads, with each thread
 - using *separate* memory arenas -- comparable to running the threads *sequentially*
 - using a *shared* memory arena with locking



2/2/08

10

measurement on the SGI Altix 200,000 states stored, 100 bytes/state



2/2/08

11

what this means...

- there is a growing performance gap
 - memory size continues to grow
 - but *cpu speed* no longer does
 - the standard approach to handling large problem sizes has stopped working
 - new algorithms, approaches are needed to leverage *large* multi-core systems
 - exploiting multi-core systems with shared memory is much harder than it would seem



2/2/08

12