



# NUMA in High-Level Languages

Patrick Siegler  
Non-Uniform Memory Architectures  
Hasso-Plattner-Institut

# Agenda

---

1. Definition of High-Level Language
2. C#
3. Java
4. Summary

# High-Level Language

---

- Interpreter, no directly machine executable format
- Platform Independence
- Automated Memory Management


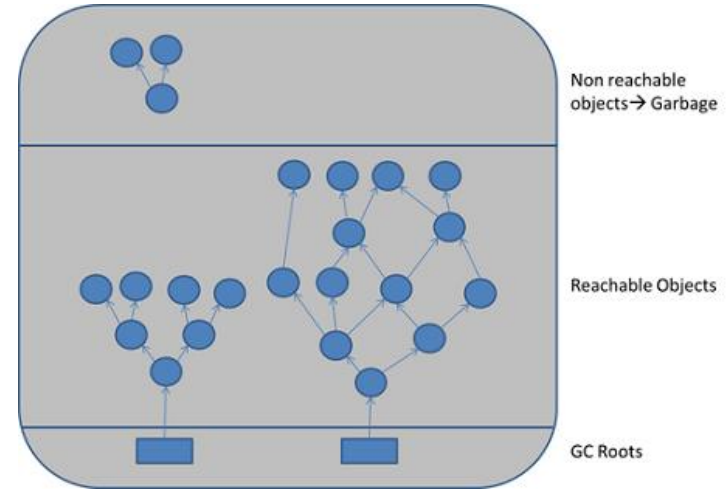
**NUMA in high-level  
languages**

Patrick Siegler  
12.12.2014

Chart 3

# GC - Short recap

- Traverse reference trees to find non-referenced objects
  - More than one GC root possible
- Reclaim space by moving referenced objects together
- Generational GC
  - many short-lived objects
  - old objects collected less frequently



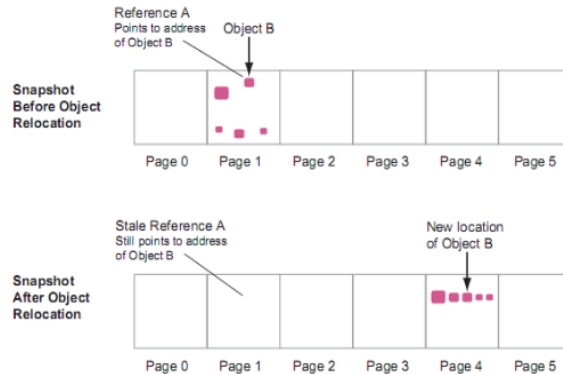
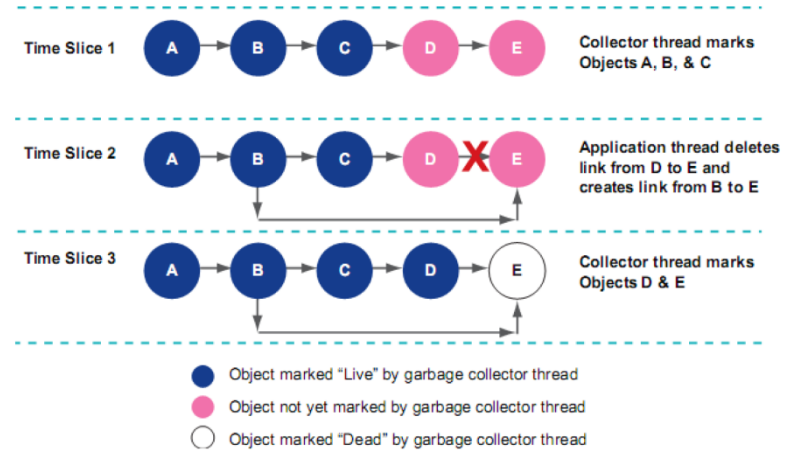
## NUMA in high-level languages

Patrick Sieglar  
12.12.2014

Chart 4

# Concurrent GC

- Difficult on multi-threaded systems
  - Modification of references during scanning
  - Lock Contention around MM data structures
  - References may be outdated
  
- Stop-the-World at some point



**NUMA in high-level languages**

Patrick Siegler  
12.12.2014

Chart 5

# GC on NUMA Systems

---

- GC compacting is copying memory
  - Expensive across nodes
- Runtime faces same problem as OS: Who is going to use which memory?
- Young objects likely to stay on node
- Abstraction conflict
  - Programs do not want to care about hardware layout
  - Association of Threads / Tasks to nodes relevant for performance

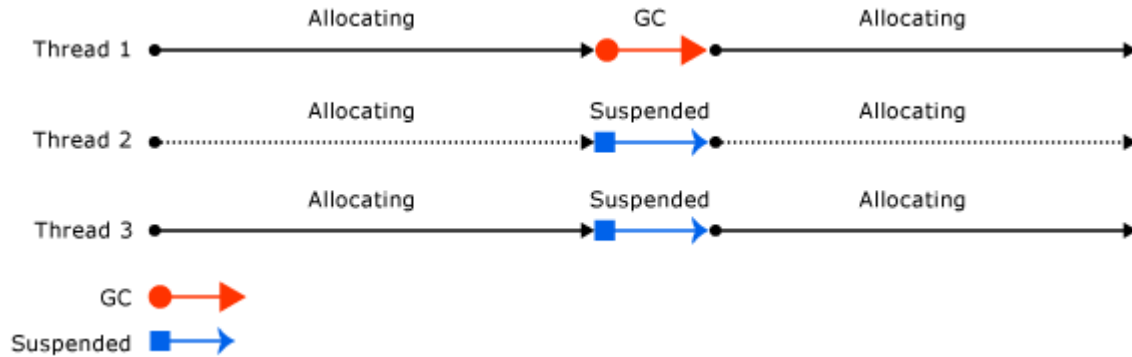
**NUMA in high-level languages**

Patrick Siegler  
12.12.2014

Chart 6

# C# - First Mutli-Processing Approach

- Stop-the-World when needed



## NUMA in high-level languages


Patrick Siegler  
12.12.2014

Chart 7

*Program, badGC*

# C# - Multi-Processing Enhancements

- Young generation collected per-thread "foreground"
- Old generation collected concurrently "background"



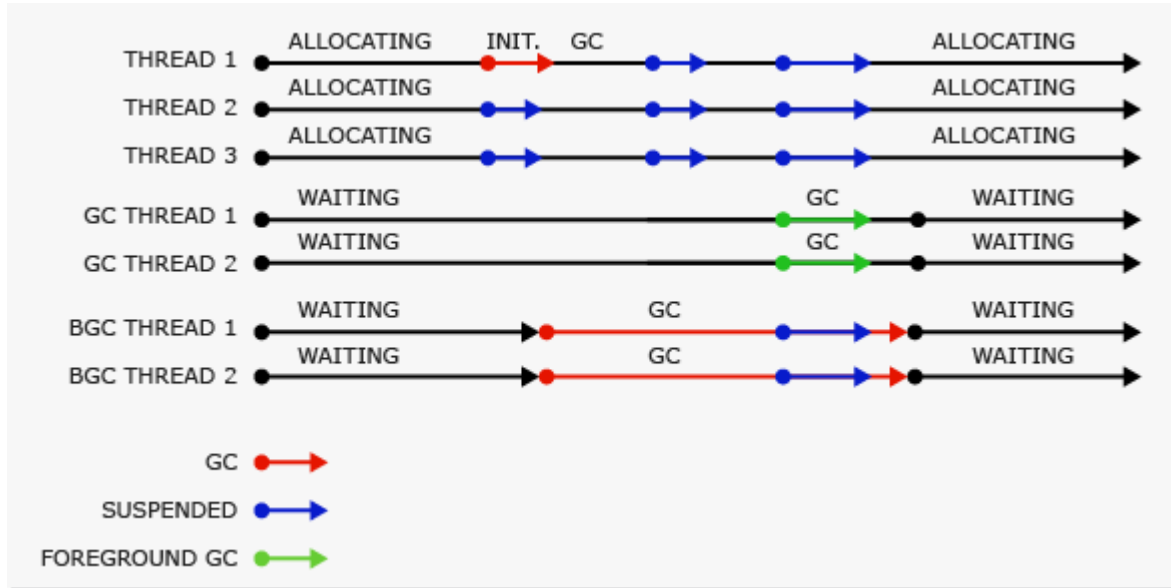
## NUMA in high-level languages

Patrick Siegler  
12.12.2014



# C# - Multi-Processing Enhancements

- “Server” GC uses dedicated high-priority threads
- 2 GC threads and a dedicated heap space per logical processor



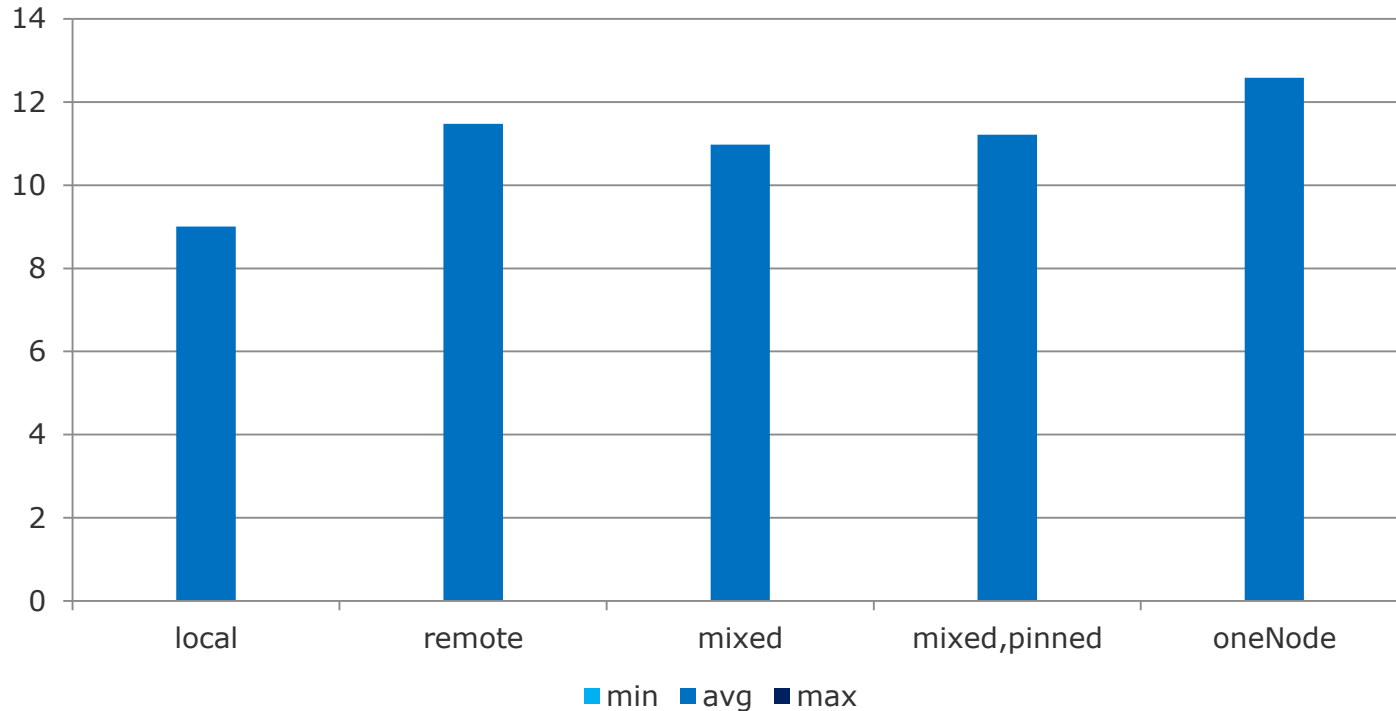
## NUMA in high-level languages

Patrick Siegler  
12.12.2014

Chart 9

*Jump, AllHelp*

# C# - Shared data access




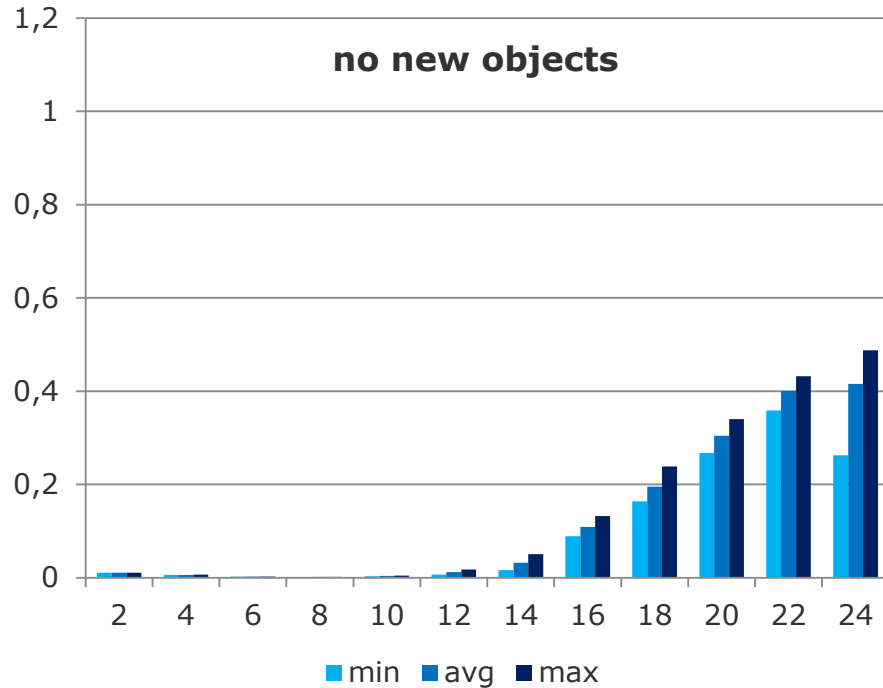
## NUMA in high-level languages

Patrick Sieglar  
12.12.2014

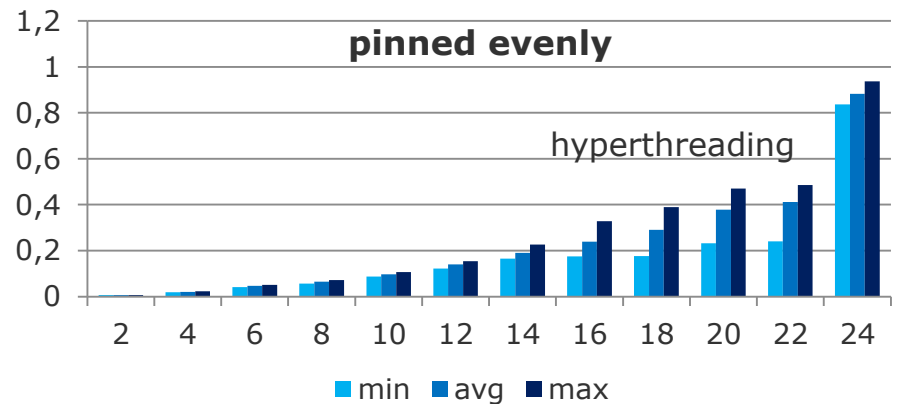

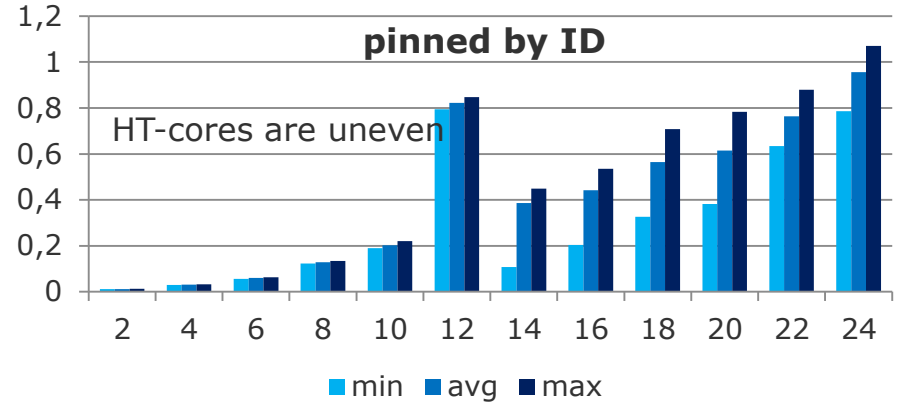
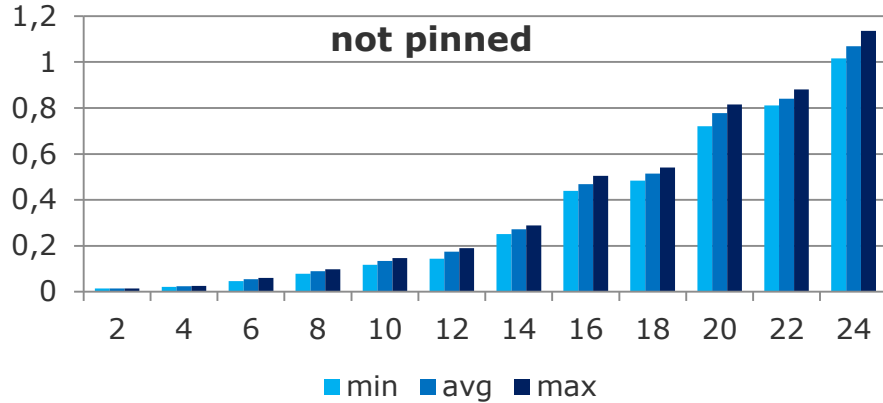
Chart 10

*MemTravels*

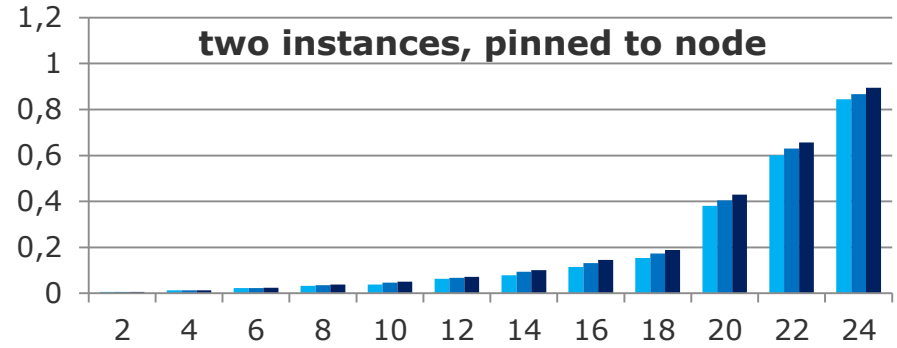
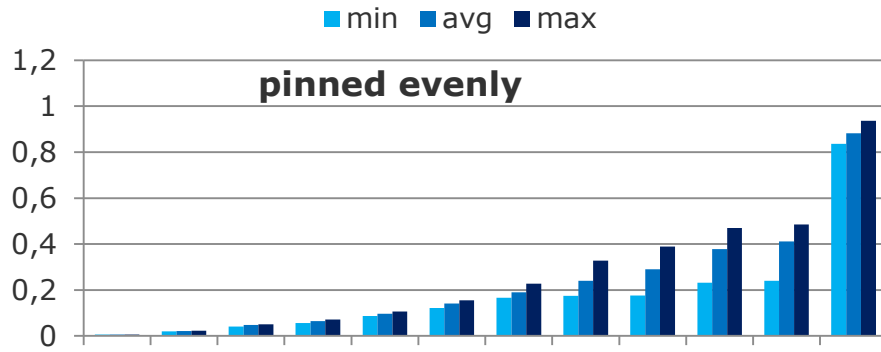
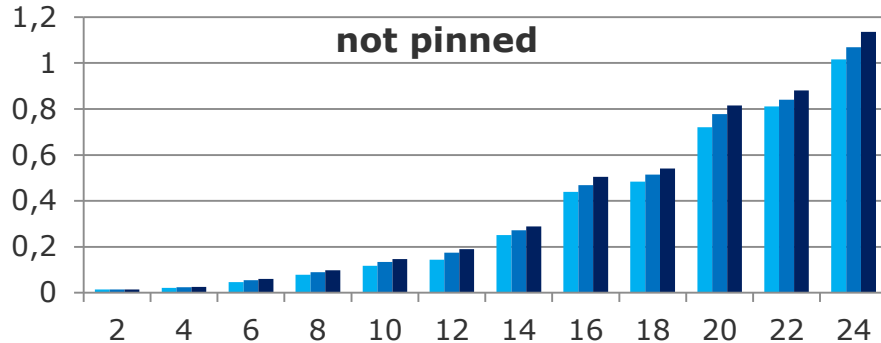
# C# - Cost of GC



# C# - Manually pinning



# C# - Single instance vs Two instances



min avg max

# C# - Summary

---

- Windows
  - Unpinned threads jump (away from their memory)
  - Natively pinned threads increase performance by >50%
  - Interconnect usage n/a on test system
  
- Linux
  - Mono's GC seems to suffer from lock contention

**NUMA in high-level languages**

Patrick Siegler  
12.12.2014

Chart **14**

# Java

---


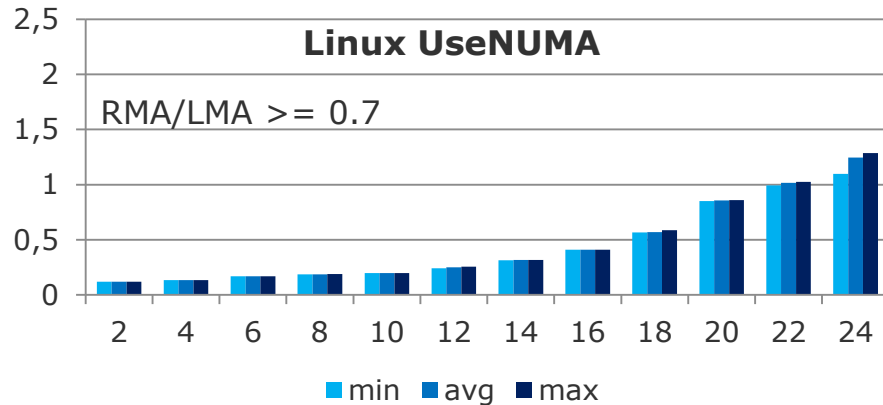
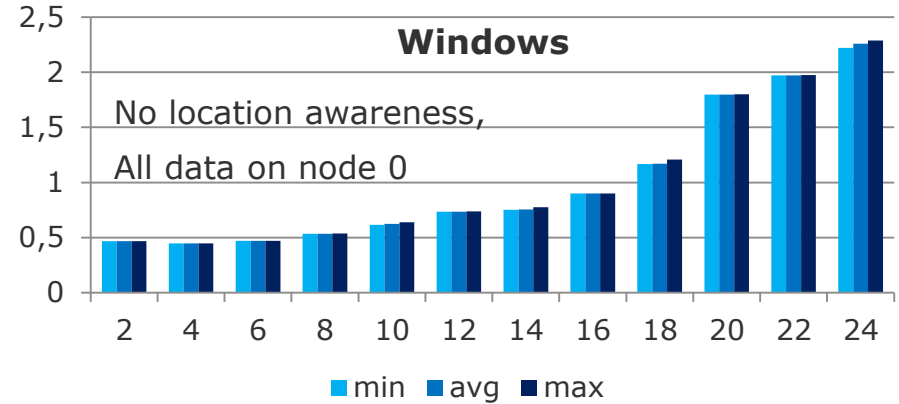
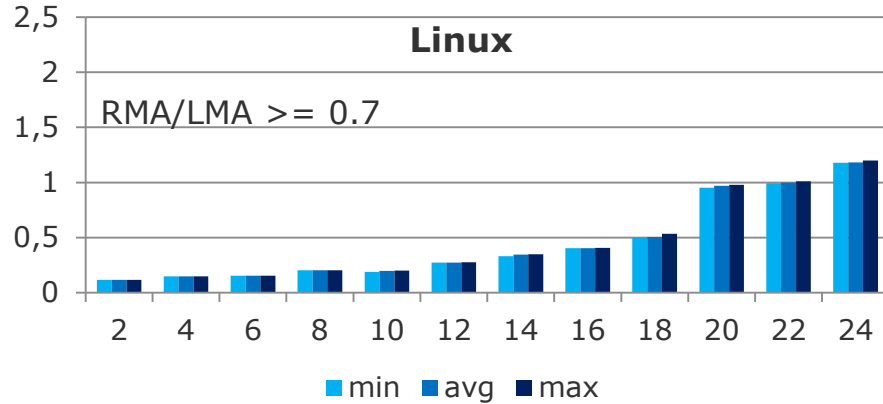
- Various virtual machines offer many GCs with varying levels of concurrency
  
- Thread-Local Allocation Buffers (TLABs)
  - synchronization-free allocation
  - no NUMA-awareness
  
- Parallel Scavenger GC (not concurrent, `-XX:+UseParallelGC`)
  - `-XX:+UseNUMA` since Java 6u2 (+40% in SPEC JBB 2005)
    - per-node regions
    - page interleaving for old and permanent generation
  - `-XX:+UseNUMAInterleaving` on Windows

## NUMA in high-level languages

Patrick Siegler  
12.12.2014

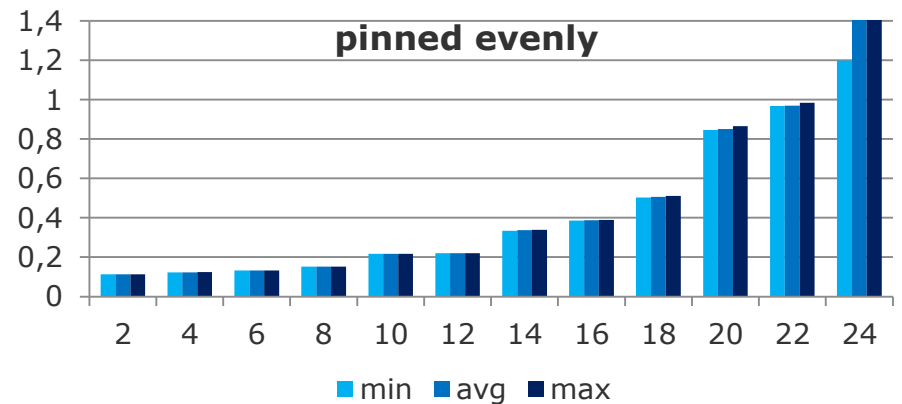
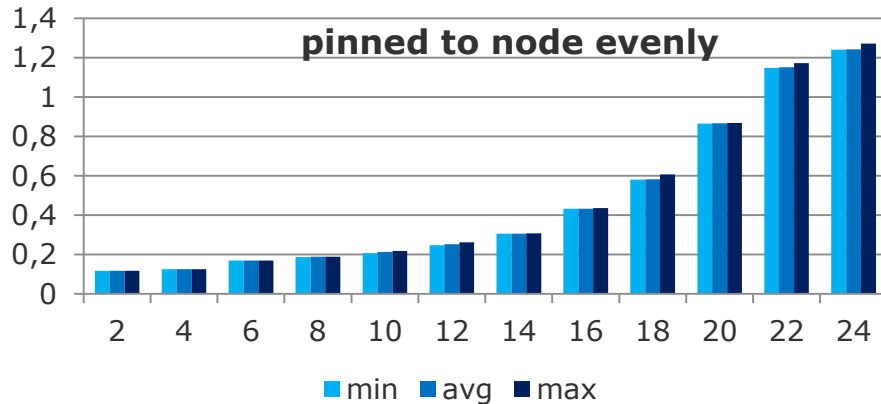
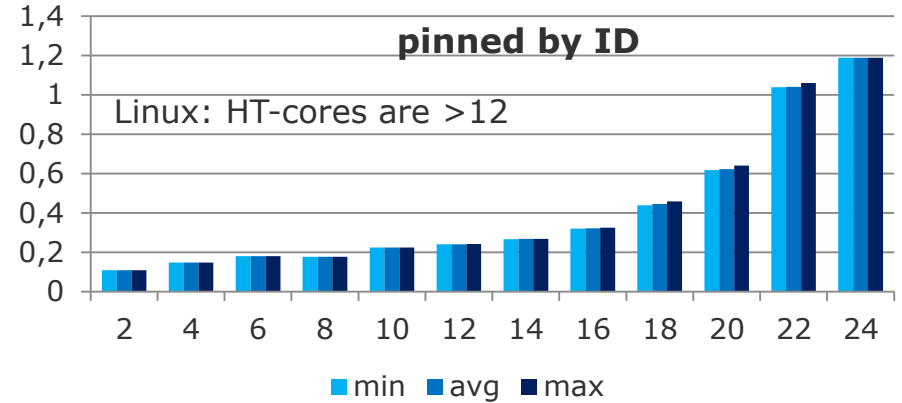
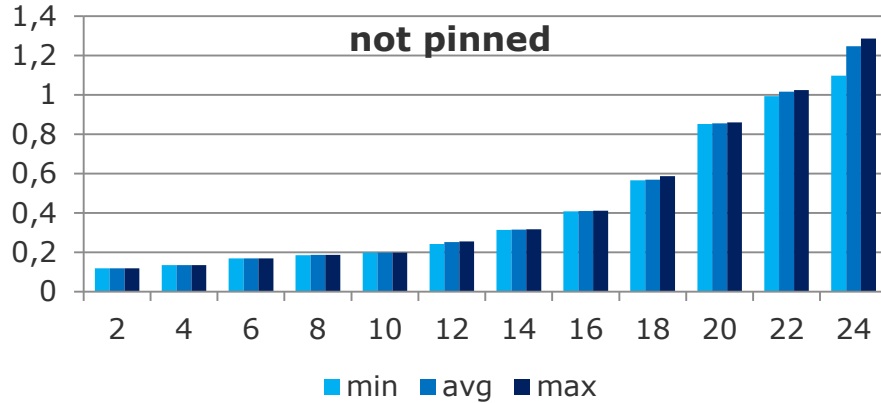
Chart **15**

# Java - OS differences (per-thread data)

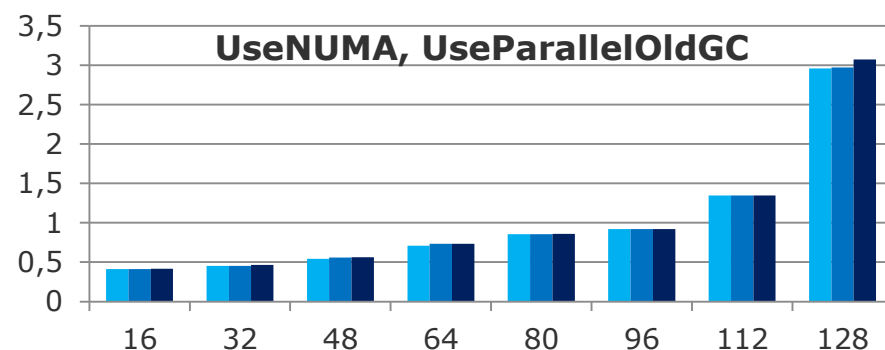
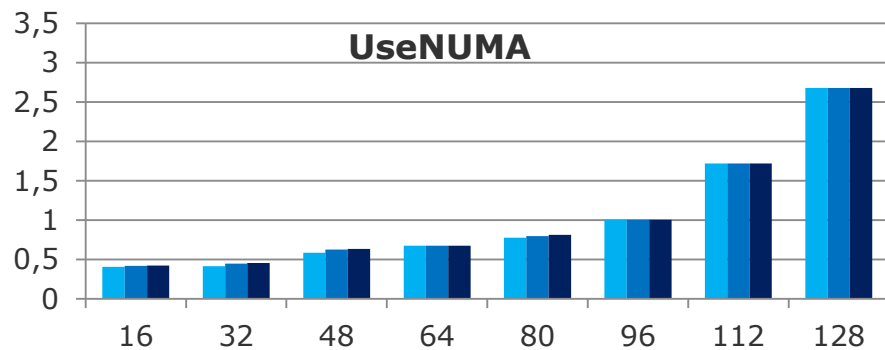
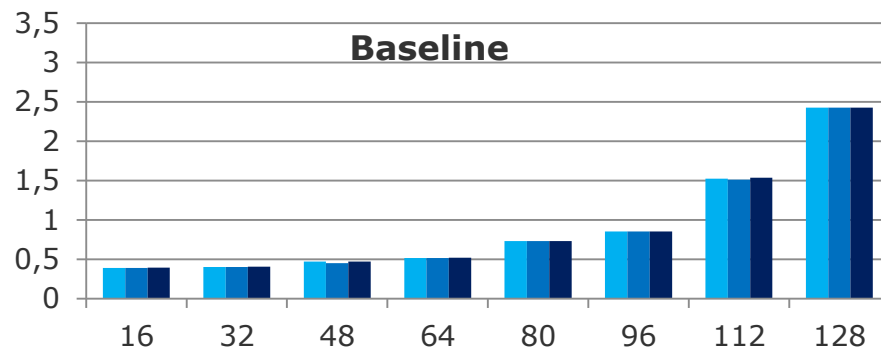




# Java - manually pinning



# Java - 8 nodes



# Java - Summary

---

- Windows
  - Limited support, only interleaving old generation
  - Unpinned threads jump
- Linux
  - Threads jump less often
  - numatop: >0.7 with UseNUMA and each thread has own data

**NUMA in high-level languages**

Patrick Siegler  
12.12.2014

Chart **19**

# Summary

---

- No explicit API in high-level languages
- Potentially catastrophic consequences of choice of GC
- Newer GC mechanisms are “NUMA-aware”
  - New allocations happen in a node-local buffer
  - Old generations are interleaved between all nodes
  - Only GC is optimized
- Varying support on different OSes
- Under-committing allows GC to operate concurrently

## **NUMA in high-level languages**

Patrick Siegler  
12.12.2014

Chart **20**

# Future and unexplored issues

---

- Improving OS schedulers will also apply to high-level languages
- Tracing and performance counters to determine which memory is used
- Actual low-level instruction flow and hyperthreading
- Actual layout of objects in memory esp. after compactation
  - Sufficient size should reduce caching effects
- Different JVM implementations

**NUMA in high-level languages**

Patrick Siegler  
12.12.2014

Chart **21**

# Sources

---

- <http://blogs.msdn.com/b/dotnet/archive/2012/07/20/the-net-framework-4-5-includes-new-garbage-collector-enhancements-for-client-and-server-apps.aspx>
- [http://msdn.microsoft.com/en-gb/library/ee787088.aspx#workstation\\_and\\_server\\_garbage\\_collection](http://msdn.microsoft.com/en-gb/library/ee787088.aspx#workstation_and_server_garbage_collection)
- <http://msdn.microsoft.com/en-us/magazine/cc163552.aspx>
  
- <http://docs.oracle.com/javase/7/docs/technotes/guides/vm/performance-enhancements-7.html>
- [www.azulsystems.com/sites/default/files//images/wp\\_pgc\\_zing\\_v5.pdf](http://www.azulsystems.com/sites/default/files//images/wp_pgc_zing_v5.pdf)
- <http://www.oracle.com/technetwork/java/javase/memorymanagement-whitepaper-150215.pdf>
- <http://developer.amd.com/community/blog/2011/10/27/a-new-numa-option-for-windows-on-the-hotspot-jvm/>
- [https://blogs.oracle.com/jonthecollector/entry/help\\_for\\_the\\_numa\\_weary](https://blogs.oracle.com/jonthecollector/entry/help_for_the_numa_weary)
- <http://javabook.compuware.com/content/memory/how-garbage-collection-works.aspx>

## **NUMA in high-level languages**

Patrick Siegler  
12.12.2014

Chart **22**