

Optimising Cloud Computing with SBSE

David R. White & Jeremy Singer
{david.r.white, jeremy.singer}@glasgow.ac.uk
University of Glasgow

Monday 25 July 2011

OUTLINE

VIRTUAL MACHINES

OPPORTUNITIES FOR SBSE

TAKE-HOME

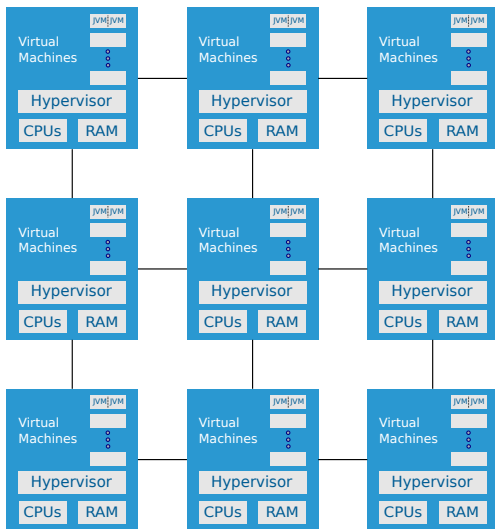
OUTLINE

VIRTUAL MACHINES

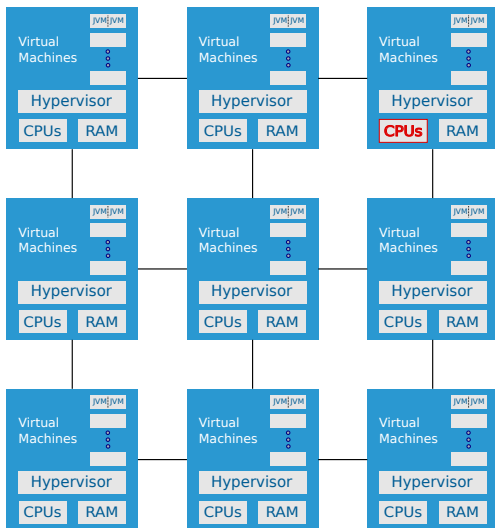
OPPORTUNITIES FOR SBSE

TAKE-HOME

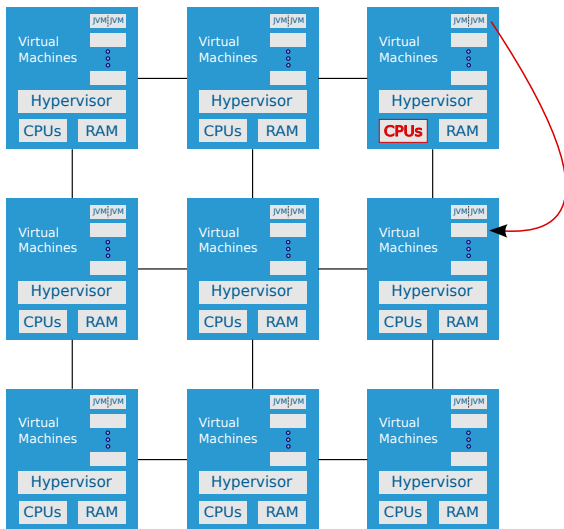
OPTIMISING VIRTUAL MACHINE MANAGEMENT



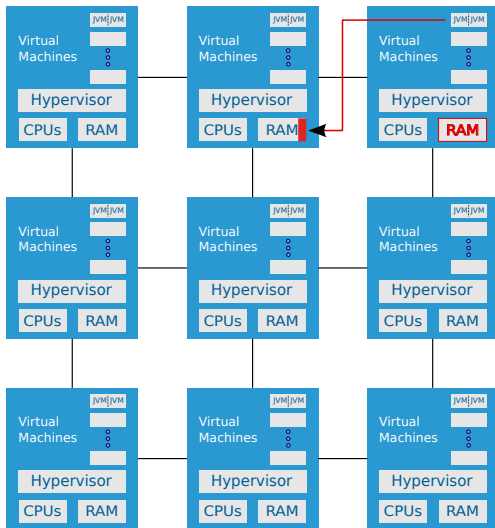
OPTIMISING VIRTUAL MACHINE MANAGEMENT



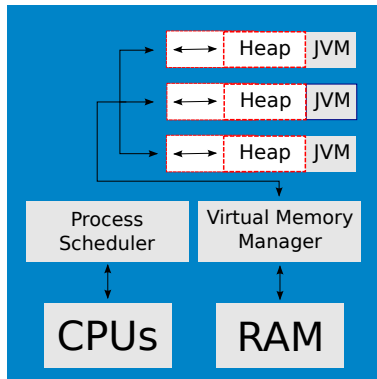
OPTIMISING VIRTUAL MACHINE MANAGEMENT



OPTIMISING VIRTUAL MACHINE MANAGEMENT



OPTIMISING A SINGLE JAVA VIRTUAL MACHINE



Competition for resources must be managed.

JVM CONFIGURATION

What heap sizes should I use?

`-Xms<number><unit>` Initial size of heap

`-Xmx<number><unit>` Maximum size of heap

JVM CONFIGURATION - II

Which Garbage Collector should I use?

- ▶ SemiSpace
- ▶ MarkSweep
- ▶ GenCopy
- ▶ GenMS
- ▶ CopyMS
- ▶ RefCount

JVM CONFIGURATION - II

Which Garbage Collector should I use?

- ▶ SemiSpace
- ▶ MarkSweep
- ▶ GenCopy
- ▶ GenMS
- ▶ CopyMS
- ▶ RefCount

Many, many more decisions to be made...

JVM CONFIGURATION: DETAIL - I

schedulingMultiplier	Should we eagerly finish sweeping at the start of a collection
eagerCompleteSweep	Should memory be protected on release?
protectOnRelease	Should finalization be disabled?
noFinalizer	Should reference type processing be disabled?
noReferenceTypes	Should a major GC be performed when a system GC is triggered?
fullHeapSystemGC	Should we ignore calls to java.lang.System.gc?
ignoreSystemGC	Should we shrink/grow the heap to adjust to application working set?
variableSizeHeap	If true, all spaces are eagerly demand zero mapped at boot time
eagerMmapSpaces	Number of bits to use for the header cycle of mark sweep spaces
markSweepMarkBits	Force a collection after this much allocation
stressFactor	Trigger a GC if the meta data volume grows to this limit
metaDataLimit	Bound the maximum size of the nursery to this value
boundedNursery	Fix the minimum and maximum size of the nursery to this value
fixedNursery	Number of GC threads to use
threads	
enable_recompilation	Should the adaptive system recompile hot methods?
enable_precompile	Should the adaptive system precompile all methods given in the advice file before the user thread is started?
adaptive_inlining	Should we use adaptive feedback-directed inlining?
osr_promotion	Should AOS promote baseline-compiled methods to opt?
background_recompilation	Should recompilation be done on a background thread or on next invocation?
method_sample_size	How many timer ticks of method samples to take before reporting method hotness to controller
decay_frequency	After how many clock ticks should we decay
dcg_decay_rate	What factor should we decay call graph edges hotness by
dcg_sample_size	After how many timer interrupts do we update the weights in the dynamic call graph?
inline_ai_seed_multiplier	Initial edge weight of call graph is set to AI_SEED_MULTIPLIER * (1/AI_CONTROL_POINT)
inline_ai_hot_callsite_threshold	What percentage of the total weight of the dcg demarcates warm/hot edges
offlinePlan	Name of offline inline plan to be read and used for inlining
early_exit_time	Value of controller clock at which AOS should exit if EARLY_EXIT is true
invocation_count_threshold	Invocation count at which a baseline compiled method should be recompiled
invocation_count_opt_level	Opt level for recompilation in invocation count based system
counter_based_sample_interval	What is the sample interval for counter-based sampling
max_opt_level	The maximum optimization level to enable.
focus_effort	Focus compilation effort based on frequency profile data
reads_kill	Should we constrain optimizations by enforcing reads-kill?
inline	Inline statically resolvable calls
inline_guarded	Guarded inlining of non-final virtual calls
inline_guarded_interfaces	Speculatively inline non-final interface calls

JVM CONFIGURATION: DETAIL - II

<code>inline_preex</code>	Pre-existence based inlining
<code>simplify_integer_ops</code>	Simplify operations on integers
<code>simplify_long_ops</code>	Simplify operations on longs
<code>simplify_float_ops</code>	Simplify operations on floats
<code>simplify_double_ops</code>	Simplify operations on floats
<code>simplify_ref_ops</code>	Simplify operations on references
<code>simplify_tib_ops</code>	Simplify operations on TIBs
<code>simplify_field_ops</code>	Simplify operations on fields
<code>simplify_chase_final_fields</code>	Chase final fields avoiding loads at runtime
<code>local_constant_prop</code>	Perform local constant propagation
<code>local_copy_prop</code>	Perform local copy propagation
<code>local_cse</code>	Perform local common subexpression elimination
<code>local_expression_folding</code>	Should we try to fold expressions with constants locally?
<code>control_static_splitting</code>	CFG splitting to create hot traces based on static heuristics
<code>control_unwhile</code>	Turn whiles into untils
<code>escape_simple_ipa</code>	Eagerly compute method summaries for simple escape analysis
<code>escape_scalar_replace_aggregates</code>	If possible turn aggregates (objects) into variable definition/uses
<code>escape_monitor_removal</code>	Try to remove unnecessary monitor operations
<code>escape_invokee_thread_local</code>	Compile the method assuming the invokee is thread-local. Cannot be properly set on command line.
<code>ssa</code>	Should SSA form be constructed on the HIR?
<code>ssa_expression_folding</code>	Should we try to fold expressions with constants in SSA form?
<code>ssa_redundant_branch_elimination</code>	Eliminate redundant conditional branches
<code>ssa_licm_ignore_pei</code>	Assume PEIs do not throw or state is not observable
<code>ssa_load_elimination</code>	Should we perform redundant load elimination during SSA pass?
<code>ssa_coalesce_after</code>	Should we coalesce move instructions after leaving SSA?
<code>ssa_loop_versioning</code>	Create copies of loops where runtime exceptions are checked prior to entry
<code>ssa_live_range_splitting</code>	Split live ranges using LIR SSA pass?
<code>ssa_gcp</code>	Perform global code placement
<code>ssa_gcse</code>	Perform global code placement
<code>ssa_global_bounds</code>	Perform (incomplete/unsafe) global Array Bound Check elimination on Demand
<code>ssa_splitblock_to_avoid_rename</code>	When leaving SSA create blocks to avoid renaming variables
<code>ssa_splitblock_for_local_live</code>	When leaving SSA create blocks for local liveness
<code>ssa_splitblock_into_infrequent</code>	When leaving SSA create blocks to avoid adding code to frequently executed blocks
<code>reorder_code</code>	Reorder basic blocks for improved locality and branch prediction
<code>reorder_code_ph</code>	Reorder basic blocks using Pettis and Hansen Algo2
<code>h2l_inline_new</code>	Inline allocation of scalars and arrays
<code>h2l_inline_write_barrier</code>	Inline write barriers for generational collectors
<code>h2l_inline_primitive_write_barrier</code>	Inline primitive write barriers for certain collectors
<code>h2l_no_callee_exceptions</code>	Assert that any callee of this compiled method will not throw exceptions. Cannot be properly set on command line.
<code>h2l_call_via_jtoc</code>	Plant virtual calls via the JTOC rather than from the tib of an object when possible
<code>l2m_handler_liveness</code>	Store liveness for handlers to improve dependence graph at PEIs
<code>l2m_schedule_prepass</code>	Perform prepass instruction scheduling

JVM CONFIGURATION: DETAIL - III

regalloc_coalesce_moves	Attempt to coalesce to eliminate register moves?
regalloc_coalesce_spills	Attempt to coalesce stack locations?
adaptive_instrumentation_sampling	Perform code transformation to sample instrumentation code.
adaptive_no_duplication	When performing inst. sampling, should it be done without duplicating code?
adaptive_processor_specific_counter	Should there be one CBS counter per processor for SMP performance?
adaptive_remove_yp_from_checking	Should yieldpoints be removed from the checking code (requires finite sample interval)
osr_guarded_inlining	Insert OSR point at off branch of guarded inlining?
osr_inline_policy	Use OSR knowledge to drive more aggressive inlining?
profile_edge_count_input_file	Input file of edge counter profile data
profile_infrequent_threshold	Cumulative threshold which defines the set of infrequent basic blocks
profile_cbs_hotness	Threshold at which a conditional branch is considered to be skewed
escape_max_array_size	Maximum size of array to replaced with registers by simple escape analysis
ssa_load_elimination_rounds	How many rounds of redundant load elimination will we attempt?
l2m_max_block_size	Maximum size of block for BURS, larger blocks are split
regalloc_simple_spill_cost_move_factor	spill penalty for move instructions
regalloc_simple_spill_cost_memory_operand_factor	spill penalty for registers used in memory operands
control_tableswitch_cutoff	If a tableswitch comprises this many or fewer comparisons convert it into multiple if-then-else style br
control_cond_move_cutoff	How many extra instructions will we insert in order to remove a conditional branch?
control_unroll_log	Unroll loops. Duplicates the loop body 2^n times.
control_static_splitting_max_cost	Upper bound on the number of instructions duplicated per block when trying to create hot traces with sta
control_well_predicted_cutoff	Don't replace branches with conditional moves if they are outside of the range of 0.5 +- this value
inline_max_target_size	Static inlining heuristic: Upper bound on callee size
inline_max_inline_depth	Static inlining heuristic: Upper bound on depth of inlining
inline_max_always_inline_target_size	Static inlining heuristic: Always inline callees of this size or smaller
inline_massive_method_size	Static inlining heuristic: If root method is already this big, then only inline trivial methods
inline_max_arg_bonus	Maximum bonus for reducing the perceived size of a method during inlining.
inline_precise_reg_array_arg_bonus	Bonus given to inlining methods that are passed a register of a known precise type.
inline_declared_aastored_array_arg_bonus	Bonus given when there's potential to optimize checkstore portion of aastore bytecode on parameter
inline_precise_reg_class_arg_bonus	Bonus given to inlining methods that are passed a register of a known precise type.
inline_extant_reg_class_arg_bonus	Bonus given to inlining methods that are passed a register that's known not to be null.
inline_int_const_arg_bonus	Bonus given to inlining methods that are passed an int constant argument
inline_null_const_arg_bonus	Bonus given to inlining methods that are passed a null constant argument
inline_object_const_arg_bonus	Bonus given to inlining methods that are passed an object constant argument
inline_call_depth_cost	As we inline deeper nested methods what cost (or bonus) do we wish to give to deter (or encourage) nesti
inline_ai_max_target_size	Adaptive inlining heuristic: Upper bound on callee size
inline_ai_min_callsite_fraction	Adaptive inlining heuristic: Minimum fraction of callsite distribution for guarded inlining of a callee

OUTLINE

VIRTUAL MACHINES

OPPORTUNITIES FOR SBSE

TAKE-HOME

WHY SBSE?

Manual optimisation is difficult: we have a complex system

Variables include software's behaviour, phase, and interactions

Solutions are non-obvious and require creativity.

A COMMERCIALY RELEVANT PROBLEM FOR GP

“I think GP has a toy problem problem.”

Sean Luke, June 2011.

This is most certainly *not* a toy problem!

EXISTING OPTIMISATIONS

How are decisions made at the moment?

EXISTING OPTIMISATIONS

How are decisions made at the moment?

Answer: they're often not - just deferred!

EXISTING OPTIMISATIONS

How are decisions made at the moment?

Answer: they're often not - just deferred!

Most recent version of Hotspot JVM has some adaptation.

EXISTING OPTIMISATIONS

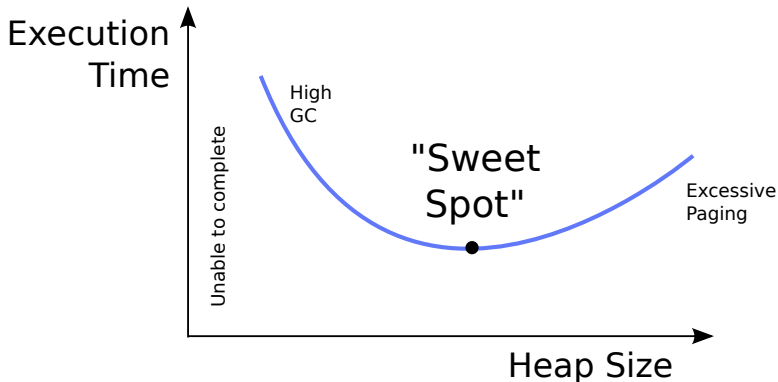
How are decisions made at the moment?

Answer: they're often not - just deferred!

Most recent version of Hotspot JVM has some adaptation.

This is a case of "Best effort" - so why not GP?

A CONCRETE EXAMPLE: HEAP SIZE CONTROL



JIKES RVM HEAP RESIZING

		Heap Occupancy					
		0.00	0.10	0.30	0.60	0.80	1.00
GC Overhead	0.00	0.90	0.90	0.95	1.00	1.00	1.00
	0.01	0.90	0.90	0.95	1.00	1.00	1.00
	0.02	0.95	0.95	1.00	1.00	1.00	1.00
	0.07	1.00	1.00	1.10	1.15	1.20	1.20
	0.15	1.00	1.00	1.20	1.25	1.35	1.30
	0.40	1.00	1.00	1.25	1.30	1.50	1.50
	1.00	1.00	1.00	1.25	1.30	1.50	1.50

Look-up table for heap-resize coefficient.

DESIGN DECISIONS

(Private Communication)

*“...back in 2003 [anon] and I did some experimental tuning and came up with the numbers by eyeballing things. At the time, it seemed to be **somewhat stable** and making reasonable decisions but that was also about 4 major versions ago and I don't think anyone has really looked at it seriously since then. I think there was **some amount of sensitivity** to the values...”*

OUTLINE

VIRTUAL MACHINES

OPPORTUNITIES FOR SBSE

TAKE-HOME

SUMMARY

- ▶ Cloud infrastructure design presents new software engineering challenges.
- ▶ Scheduling and memory management are open to optimisation.
- ▶ Many problems are amenable to SBSE.
- ▶ We are looking for collaborators!

...and relevant existing work!

{david.r.white, jeremy.singer}@glasgow.ac.uk

SEE ALSO

Cloud computing: state-of-the-art and research challenges.

Zhang et al.

Journal of Internet Services and Applications, 1/1, 2010. pp. 7-18.

Overdriver: handling memory overload in an oversubscribed cloud.

Williams et. al.

Proceedings of VEE 2011.

Previous work on service (application) level by Wada et al. and

Nallur et al.

–

{david.r.white, jeremy.singer}@glasgow.ac.uk