An Efficient Sensitivity Analysis Method for Large Cloud Simulations

July 8, 2011 IEEE Cloud 2011

K. Mills, J. Filliben and C. Dabrowski
NIST

image generated with http://www.wordle.net/ applied to source code for Koala Cloud Simulator created by Mills and Dabrowski

We Developed a 4-Step Method* to Compare Resource Allocation Algorithms in Large Distributed Systems

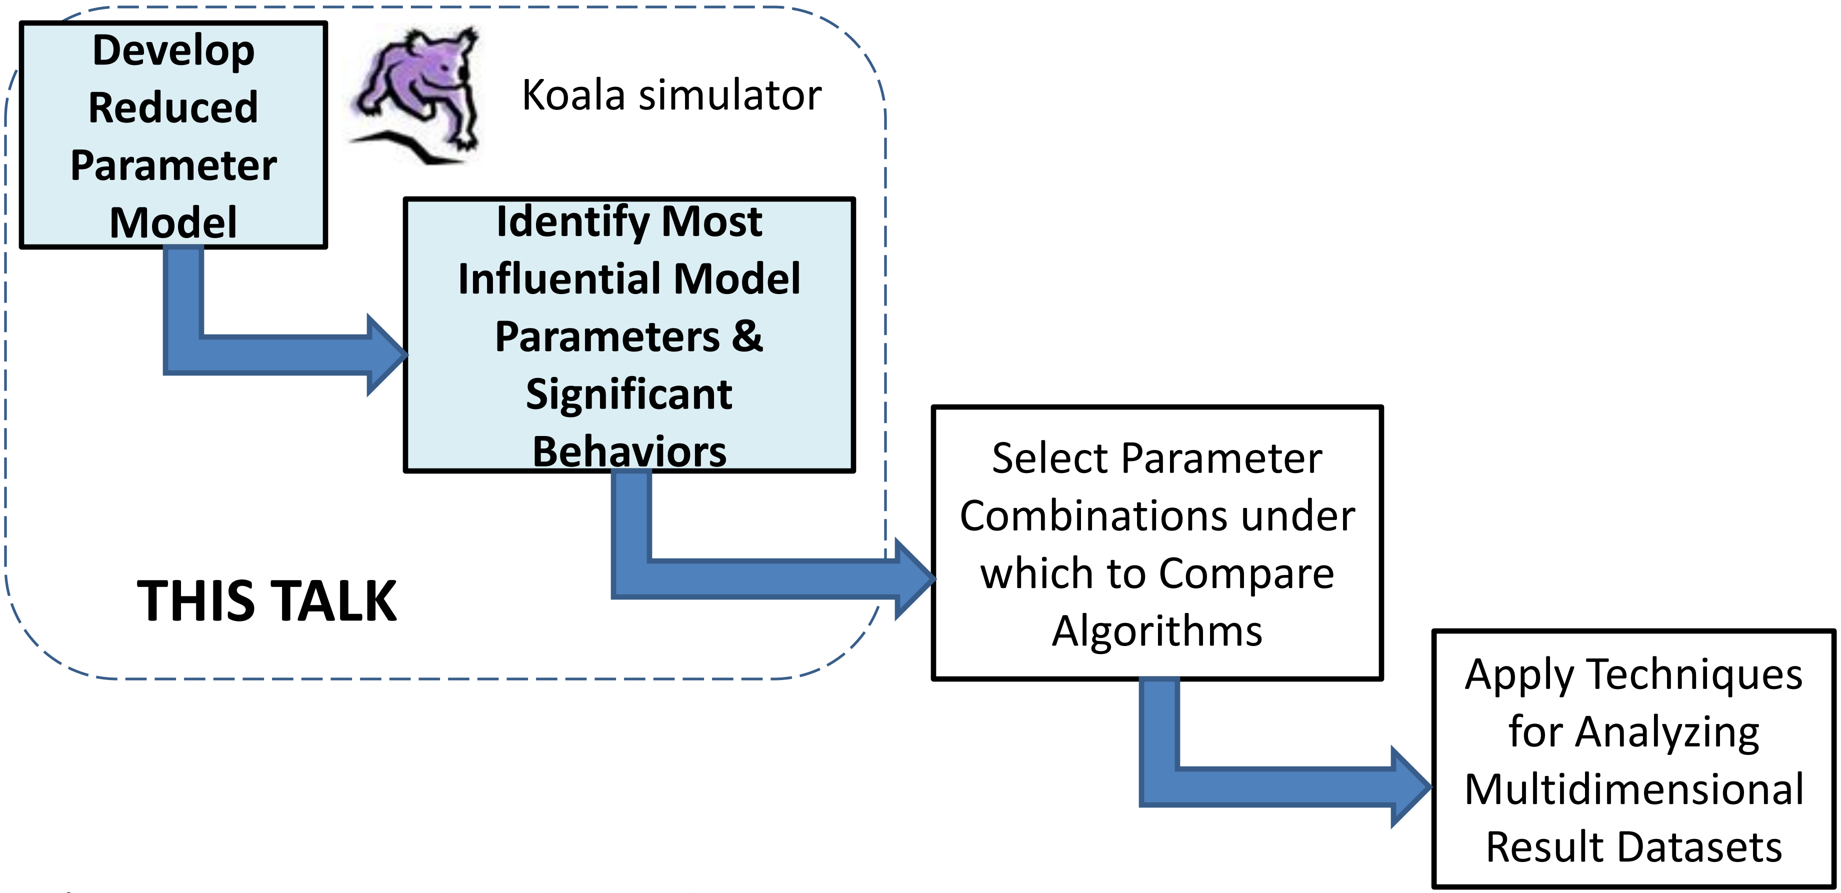

*Previously, we applied this method to compare congestion-control algorithms proposed for the Internet

Talk Synopsis

Problem: Simulations of large cloud systems typically require hundreds of parameters that can take on billions of values and that can also report hundreds of response variables.

$$\underbrace{y_1, \ldots, y_z}_{\text{Response State-Space}} = f(\underbrace{x_{1|[1,\ldots,\ell]} \ldots, x_{p|[1,\ldots,\ell]}}_{\text{Stimulus State-Space}})$$

How can one identify the most significant parameters to simulate and responses to analyze?

We base our study on the ***Koala***  infrastructure cloud simulator

$(2^{32})^{82} \rightarrow O(10^{789})$ [10^{80} = atoms in visible universe]

Group related remaining parameters – reduce by 59 parameters

$(2^{32})^{23} \rightarrow O(10^{221})$

Fix parameters not considered germane – reduce by 12 parameters

$(2^{32})^{11} \rightarrow O(10^{105})$

Select only 2 values for each parameter

$2^{11} \rightarrow 2048$

Use experiment design theory to reduce parameter combinations to 256

$2^{11-5} \rightarrow 64$

Use sensitivity analysis results to identity six most significant parameters

$2^{6-1} \rightarrow 32$

Use experiment design theory again to reduce parameter combinations to 32

Model Reduction

Level Reduction

Sensitivity Analysis

Experiment Design Theory

Outline

- Overview of ***Koala*** Infrastructure Cloud Simulator
- Sensitivity Analysis Experiment Design
- Reduction of Response Dimensionality
- Identification of Significant Parameters
- Findings and Ongoing Work

Overview of *Koala* Infrastructure Cloud Simulator

Schematic of *Koala* IaaS Cloud Computing Model

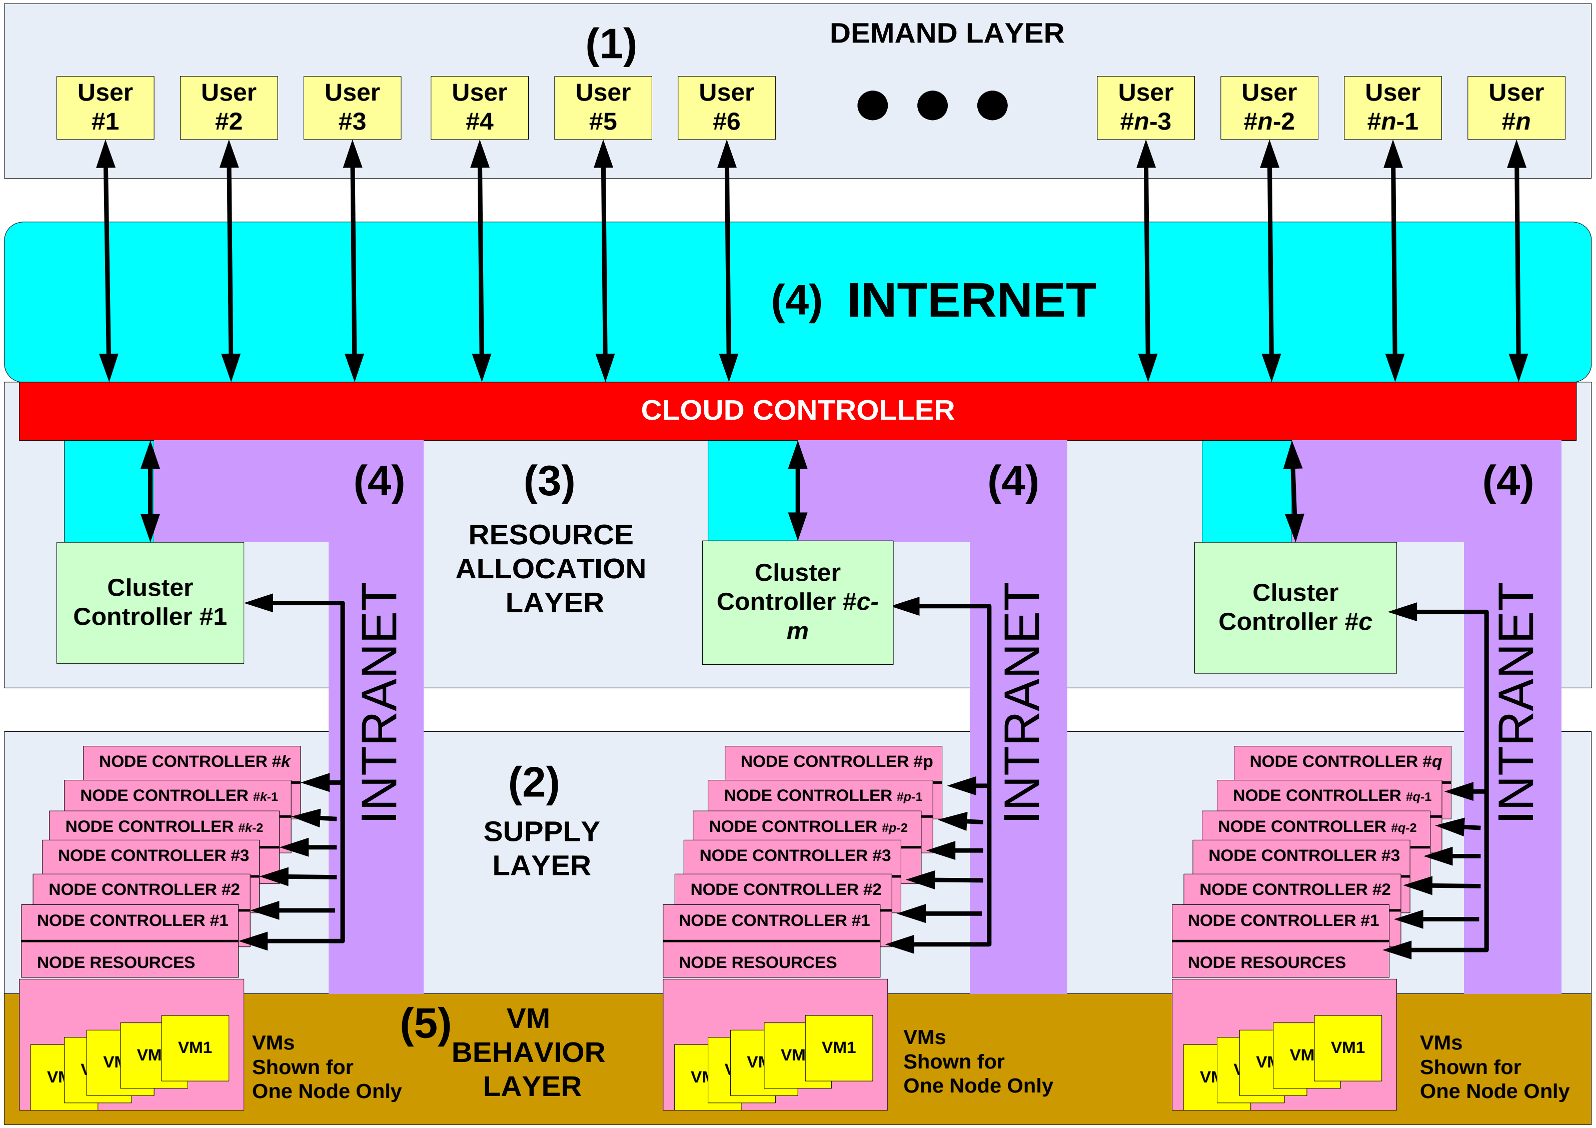

Virtual Machine (VM) Types* Simulated in *Koala*

VM Types are offered by the Cloud provider and requested by Cloud users

VM Type	Virtual Cores #	Virtual Cores Speed (GHz)	Virtual Block Devices #	Virtual Block Devices Size (GB) of Each	# Virtual Network Interfaces	Memory (GB)	Instruct. Arch.
M1 small	1	1.7	1	160	1	2	32-bit
M1 large	2	2	2	420	2	8	64-bit
M1 xlarge	4	2	4	420	2	16	64-bit
C1 medium	2	2.4	1	340	1	2	32-bit
C1 xlarge	8	2.4	4	420	2	8	64-bit
M2 xlarge	8	3	1	840	2	32	64-bit
M4 xlarge	8	3	2	850	2	64	64-bit

*Inspired by Amazon Elastic Compute Cloud VM Types

Description of User Types Simulated in *Koala*

We created different classes of demand, such as processing users (PU), distributed simulation users (MS), peer-to-peer users (PS), Web service users (WS) and data search users (DS)

User Type	VM Type(s)	Max-Min VMs	Max-Max VMs
PU1	M1 small	10	100
PU3		100	500
PU5		500	1000
PU2	M1 large	10	100
PU4		100	500
PU6		500	1000
MS1	M1 xlarge	10	100
MS3	M1 xlarge	100	500

User Type	VM Type(s)	Max-Min VMs	Max-Max VMs
PS1	C1 medium	3	10
PS2		10	50
PS3		50	100
WS1	M1 large M2 xlarge C1 xlarge	1	3
WS2	M1 large M2 xlarge C1 xlarge	3	9
WS3	M1 large M2 xlarge C1 xlarge	9	12
DS1	M4 xlarge	10	100
DS2		100	500
DS3		500	1000

Finite-State Machine of Simulated User Behavior in *Koala*

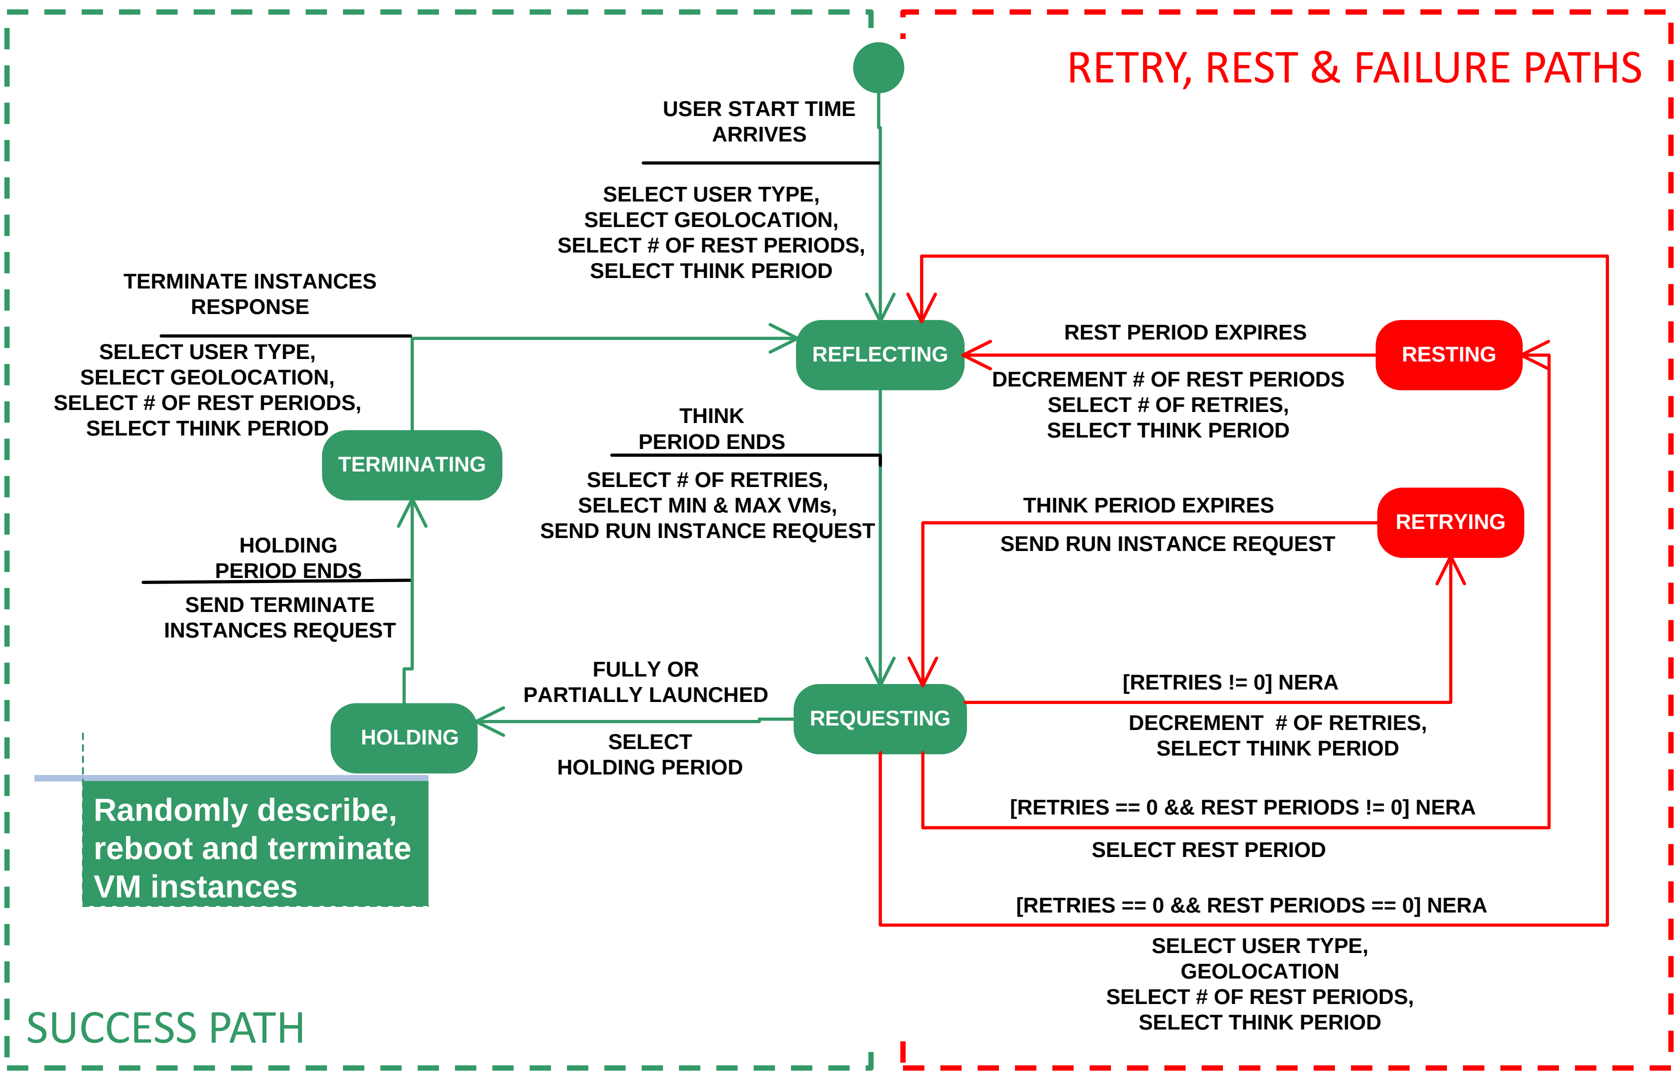

Description of Selected Platform Types Simulated in *Koala*

We created 22 platform classes, inspired by
a visit to an Amazon EC2 data center

Platform Type	Physical Cores		Memory (GB)	# Physical Disks by Size				# Network Interfaces	Instruct. Arch.
	#	Speed (GHz)		250 GB	500 GB	750 GB	1000 GB		
C2	1	1.7	16	3	0	0	0	1	32-bit
C4	1	2	16	3	0	0	0	1	32-bit
C6	2	2.4	16	0	3	0	0	1	32-bit
C8	2	2.4	32	0	3	0	0	1	64-bit
C10	4	2.4	32	0	4	0	3	1	64-bit
C12	4	2	64	0	4	0	3	2	64-bit
C14	4	3	64	0	4	0	3	2	64-bit
C16	8	3	64	0	0	4	3	2	64-bit
C18	8	3	128	0	0	4	3	4	64-bit
C20	16	3	128	0	0	0	7	4	64-bit
C22	16	3	256	0	0	0	7	4	64-bit

Sensitivity Analysis Experiment Design

Input Parameters used for Sensitivity Analysis of *Koala*

We identified 11 parameters we expected to significantly influence *Koala* behavior

Category	ID	Parameter Name
Duration	*x1*	Simulation duration in hours
Demand Layer	*x2*	Number of users
	x3	Probability of user's type
	x4	Average (and shape of) user holding time
Supply Layer	*x5*	Number of clusters
	x6	Number of nodes per cluster
	x7	Probability of platform configuration type
Resource Control Layer	*x8*	Algorithm for selecting cluster
	x9	Algorithm for selecting node
Internet/ Intranet Layer	*x10*	Number of sites for cloud components
	x11	Probability range of packet losses

Response Variables used for Sensitivity Analysis of *Koala*

We selected 40 variables that we expected to represent significant *Koala* dynamics

Category	ID	Response Name (Definition)
User-Level Responses	*y1*	User Request Rate (Requests by All Users / # User Cycles)
	y2	NERA Rate (NERAs / Requests by All Users)
	y3	Full Grant Rate (Full Grants / (Full Grants + Partial Grants))
	y4	User Arrival Rate (# User Cycles / Simulated Hours)
	y5	User Give-up Rate (# Users that Gave Up / # User Cycles)
	y6	Grant Latency (Weighted Avg. Delay in Granting VMs to Users that Got VMs)
Cloud-Level Responses	*y7*	Reallocation Rate (# Times Alternate Cluster Chosen / Requests Granted)
	y8	Full Grant Proportion (Avg. Fraction Clusters Offering Full Grants)
	y9	NERA Proportion (Avg. Fraction Clusters Reporting NERA)
	y10	vCore Utilization (Avg. Fraction of Virtual Cores Used in Cloud)
	y11	Memory Utilization (Avg. Fraction of Memory in Use in Cloud)
	y12	Disk Space Utilization (Avg. Fraction of Disk Space in Use in Cloud)
	y13	pCore Load (Avg. Virtual Cores Allocated / Physical Cores in Cloud)
	y14	Disk Count Load (Avg. Virtual Disks Allocated / Physical Disks in Cloud)
	y15	NIC Count Load (Avg. Virtual NICs Allocated / Physical NICs in Cloud)
Cluster-Level Responses	*y16*	vCore Util. Var. (Avg. Variance in vCore Utilization across Clusters)
	y17	Memory Util. Var. (Avg. Variance in Memory Utilization across Clusters)
	y18	Disk Space Util. Var. (Avg. Variance in Disk Space Utilization across Clusters)
	y19	pCore Load Var. (Avg. Variance in pCore Load across Clusters)
	y20	Disk Count Var. (Avg. Variance in Disk Count Load across Clusters)
	y21	NIC Count Var. (Avg. Variance in NIC Count Load across Clusters)
	y22	Node Reallocation Rate (# Times Alternate Node Chosen / VMs Allocated)
	y23	Cluster NERA Rate (# NERAs / # Responses Avg. across Clusters)
	y24	Cluster Full-Grant Rate (# Full Grants / # Responses Avg. across Clusters)
	y25	Allocation Rate (Times Cluster chosen / Cluster offered Avg. across Clusters)
	y26	SD-NERA (Stand. Dev. in Avg. NERA Rate across Clusters)
	y27	SD-Full-Grant (Stand. Dev. in Avg. Full-Grant Rate across Clusters)
	y28	SD-Allocation-Rate (Stand. Dev. in Allocation Rate across Clusters)
VM-Level Responses	*y29*	Current Instances (Avg. # VM Instances Extant in Cloud)
	y30	M1small Instances (Fraction of Current Instances that are M1 small VMs)
	y31	M1large Instances (Fraction of Current Instances that are M1 large VMs)
	y32	M1xlarge Instances (Fraction of Current Instances that are M1 xlarge VMs)
	y33	C1medium Instances (Fraction of Current Instances that are C1 medium VMs)
	y34	C1xlarge Instances (Fraction of Current Instances that are C1 xlarge VMs)
	y35	M2xlarge Instances (Fraction of Current Instances that are M2 xlarge VMs)
	y36	M4xlarge Instances (Fraction of Current Instances that are M4 xlarge VMs)
Message-Level Responses	*y37*	WS Message Rate (Avg. # WS Messages Send Per Simulated Hour)
	y38	Intra-Site Messages (# WS Messages Sent with Sites / # WS Messages Sent)
	y39	Inter-Site Loss Rate (Avg. Fraction of Inter-Site WS Messages Undelivered)
	y40	Intra-Site Loss Rate (Avg. Fraction of Intra-Site WS Messages Undelivered)

2-Level OFF Experiment Designs Reduce # of Parameter Combinations, While Improving Global Coverage and Minimizing Error in Effect Estimates in comparison with comparable Factor-at-a-Time (FAT) Designs

We selected two pairs of level settings (SA1 & SA2) and two system sizes (small & large)

Adopted 2-Level (2^{11-5}) "Resolution IV" OFF experiment design, requiring 64 simulations per experiment

Instantiated 4 designs, and simulated 6 repetitions (different random number seeds) with the 2 smaller designs

Required $(6 \times 2 + 2) \times 64 = 896$ simulations

	SA1-small and SA1-large		SA2-small and SA2-large	
Parameter	Plus Level	Minus Level	Plus Level	Minus Level
x1	1200 hours	600 hours	1600 hours	200 hours
x2	500 (SA1-small) 5000 (SA1-large)	250 (SA1-small) 2500 (SA1-large)	750 (SA2-small) 7500 (SA2-large)	125 (SA2-small) 1250 (SA2-large)
x3	PU1 = 0.2 PU2 = 0.2 PU3 = 0.1 PU4 = 0.1 WS1 = 0.15 WS2 = 0.07 WS3 = 0.03 PS1 = 0.1 PS2 = 0.01 MS1 = 0.1 MS3 =0.01 DS1 = 0.10 DS2 = 0.01	PU1 = 1/6 PU2 = 1/6, WS1 = 1/6 MS1 = 1/6 PS1 = 1/6 DS1 = 1/6	PU1 = 0.4 PU2 = 0.4 PU3 = 0.1 PU4 = 0.05 PU5 = 0.025 PU6 = 0.025	WS1 = 0.25 WS2 = 0.15 WS3 = 0.1 PS1 = 0.35 PS2 = 0.04 PS3 = 0.01 DS1 = 0.08 DS2 = 0.015 DS3 = 0.005
x4	8 hours ($\alpha = 1.2$)	4 hours ($\alpha = 1.2$)	12 hours ($\alpha = 1.2$)	2 hours ($\alpha = 1.2$)
x5	20 (SA1-small) 40 (SA1-large)	10 (SA1-small) 20 (SA1-large)	30 (SA2-small) 40 (SA2-large)	5 (SA2-small) 10 (SA2-large)
x6	200 (SA1-small) 1000 (SA1-large)	100 (SA1-small) 500 (SA1-large)	400 (SA2-small) 1500 (SA2-large)	50 (SA2-small) 250 (SA2-large)
x7	C22 = 1.0	C8 = 0.25 C14 = 0.25 C18 = 0.25 C22 = 0.25	C14 = 0.2 C16 = 0.2 C18 = 0.2 C20 = 0.2 C22 = 0.2	C2 = 0.1 C4 = 0.1 C6 = 0.1 C8 = 0.1 C10 = 0.1 C12 = 0.1 C16 = 0.1 C22 = 0.3
x8	Percent Allocated	Least-Full First	Percent Allocated	Least-Full First
x9	Next-Fit	First-Fit	Next-Fit	First-Fit
x10	4	1	8	1
x11	10^{-3} to 10^{-8}	10^{-4} to 10^{-9}	10^{-2} to 10^{-7}	10^{-5} to 10^{-10}

Reduction of Response Dimensionality

Correlation Analysis & Clustering (CAC) Reduces Dimensionality

We identified an **8-dimensional response space** within the 40 responses

Compute correlation coefficient (*r*) for all response pairs

Examine frequency distribution for all |*r*| to determine threshold for correlation pairs to retain; |*r*| > 0.65, here

Create clusters of mutually correlated pairs; each cluster represents one dimension

Select one response from each cluster to represent the dimension; we selected response with largest mean correlation that was not in another cluster*

Response Dimension	SA1-small (9 dimensions)	SA1-large (8 dimensions)	SA2-small (10 dimensions)	SA2-large (9 dimensions)
Cloud-wide Demand/Supply Ratio	*y1, y2,* ***y3****, y5, y6, y8, y9, y10, y13, y23, y24, y25, y29, y30, y32, y34, y36, y38*	*y1, y2,* ***y3****, y5, y6, y7, y8, y9, y10, y13, y23, y34, y25, y29, y30, y32, y33, y34, y36, y38*	*y1,* ***y2****, y3, y5, y6, y8, y9, y10, y11, y13, y14, y15, y23, y24, y25, y38*	*y1, y2, y3, y5, y6, y8, y9,* ***y23****, y24, y25, y38*
Cloud-wide Resource Usage	*y10, y11, y12, y13, y14,* ***y15***	*y10, y11, y12, y13, y14,* ***y15***	***y10****, y11, y12, y13, y14, y15*	***y10****, y11, y12, y13, y14, y15*
Variance in Cluster Load	*y16, y17, y18, y19,y20, y21,* ***y26****, y27*	*y16, y17, y18, y19,y20, y21,* ***y26****, y27*	*y16, y18, y19, y20, y21, y26,* ***y27*** ***y17*** *(Mem. Util)*	*y16, y17, y18,* ***y19****,y20, y21, y26, y27*
Mix of VM Types	*y34,* ***y35*** (WS) ***y31*** (MS)	***y31*** (MS)	*y12, y14, y15, y30, y31, y33, y34, y35,* ***y36***	*y14, y15, y30,* ***y31****, y33, y34, y35* *y15,* ***y36*** (DS)
Number of VMs	*y29,* ***y37***	***y37***	*y29,* ***y37***	***y29***
User Arrival Rate	***y4***	***y4***	***y4***	***y4****, y37*
Reallocation Rate	***y7****, y22*	*y7,* ***y22***	***y7*** (cluster) ***y22*** (node)	*y7,* ***y22***
Variance in Choice of Cluster	***y28***	***y28***	***y28***	***y28***

*Not possible for cloud-wide resource usage in SA2-small, so we selected response with highest mean correlation.

Analysis of Correlation Directionality Aids Model Verification

We checked positive and inverse correlations for reasonableness

	y1	y2	y3	y4	y5	y6	y7	y8	y9	y10	y11	y12	y13	y14	y15	y16	y17	y18	y19	y20	y21	y22	y23	y24	y25	y26	y27	y28	y29	y30	y31	y32	y33	y34	y35	y36	y37	y38
y1																																						
y2	P																																					
y3	N	N																																				
y4																																						
y5	P	P																																				
y6	P	P	N		P																																	
y7																																						
y8		N	P																																			
y9	P	P	N			P		N																														
y10			N					N	P																													
y11																																						
y12								N			P																											
y13			N						P	P																												
y14											P	P																										
y15										P		P	P	P																								
y16																																						
y17																P																						
y18																P	P																					
y19																P	P	P																				
y20																P	P	P	P																			
y21																P	P	P	P	P																		
y22							P																															
y23	P	P	N			P		N	P	P			P																									
y24		N	P					P	N	N			N										P															
y25	P	P	N		P	P		N	P														P	N														
y26		N	P													P	P	P	P	P	P																	
y27																P	P	P	P	P	P					P												
y28																																						
y29			P																						N													
y30		P	N					N	P														P	N														
y31																																						
y32		N						P	N														N	P						N								
y33	P		N			P																			P													
y34																																N						
y35																																N		P				
y36			P					P	N														N	P						N	P	P			N			
y37																													P									
y38	N	N	P		N	N		P	N														N	P	N								N					

Correlation Half-Matrix from SA1-small, which had largest number (126) of correlation pairs with |*r*| > 0.65

P = positive correlation **N** = negative correlation

Identification of Significant Parameters

Main Effects Analysis (MEA) Identifies Significant Influence of Input Parameters on Response Variables

We applied MEA to response variables selected using CAC –
this example is ***y15*** (NIC Count Load) for **experiment SA1-small**

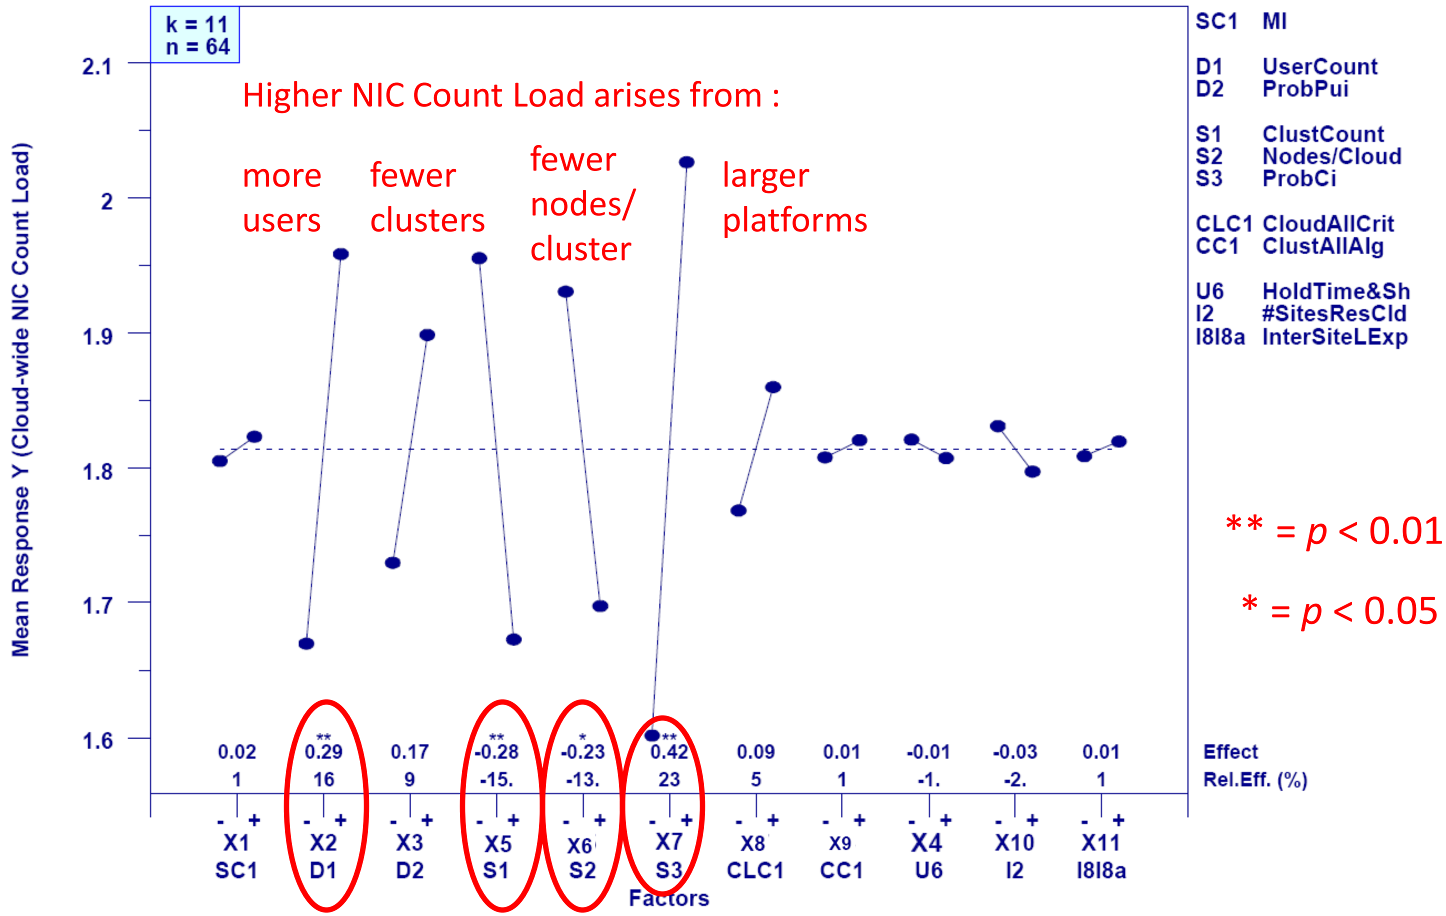

** = $p < 0.01$

* = $p < 0.05$

Most significant parameters determined through MEA of the responses selected using CAC

We computed percent of responses influenced (Ψ) for each parameter, weighting $p < 0.05$ at ½ and $p < 0.01$ at 1:

$$\Psi = (|\{y \mid p < 0.01\}| + ½\,|\{y \mid p < 0.05\}|)\,/\,|\{y\}| \times 100$$

Computed average Ψ for each parameter, weighting experiment Ψ by number of repetitions

		Input Parameter										
Experiment	**Weight**	***x1***	***x2***	***x3***	***x4***	***x5***	***x6***	***x7***	***x8***	***x9***	***x10***	***x11***
SA1 small	**6/14**	**1**	**57**	**22**	**11**	**44**	**29**	**30**	**12**	**0**	**1**	**0**
SA1 large	**1/14**	**0**	**69**	**13**	**25**	**44**	**56**	**31**	**25**	**0**	**13**	**0**
SA2 small	**6/14**	**2**	**73**	**38**	**10**	**45**	**62**	**10**	**17**	**1**	**0**	**0**
SA2 large	**1/14**	**0**	**56**	**50**	**11**	**39**	**56**	**6**	**11**	**0**	**0**	**0**
Avg. Ψ	**Est.**	**1**	**65**	**30**	**12**	**44**	**47**	**20**	**15**	**0**	**1**	**0**

green = major influence; yellow = modest influence; orange = minor influence; gray = no influence

Most significant parameters: *x2* (# users), *x5* (# clusters), and *x6* (# nodes/cluster)

Moderately influential parameters: *x3* (user types) and *x7* (platform types)

Somewhat influential parameters: *x4* (user hold time) and *x8* (cluster-selection algorithm)

No influence : *x1* (measurement interval), *x9* (node-selection algorithm), *x10* (geo-distribution of cloud components), and *x11* (packet loss prob.)

Checking Relative Effects from MEAs aids Model Verification

We averaged relative effect (Δ) over all experiment repetitions to determine how parameter increase influences direction and magnitude of effect for 8 dimensions

Dimension	Selected Response	Input Parameter										
		x1	*x2*	*x3*	*x4*	*x5*	*x6*	*x7*	*x8*	*x9*	*x10*	*x11*
Cloud-wide Demand/ Supply Ratio	*y3*	1	-38	-21	-5	37	40	25	-2	1	5	-1
Cloud-wide Resource Usage	*y15*	1	23	53	1	-22	-18	19	-1	1	-1	1
Variance in Cluster Load	*y26*	0	-101	28	-5	96	59	66	42	0	50	0
Mix of VM Types	*y31*	-1	-9	43	-3	7	9	-1	8	2	-4	0
Number of VMs	*y37*	-5	48	11	-23	79	53	-5	8	-1	4	-2
User Arrival Rate	*y4*	-17	87	2	-80	29	31	15	-4	-1	-3	-5
Reallocation Rate	*y7*	0	0	0	0	0	0	0	0	0	0	0
Variance in Cluster Choice	*y28*	6	-12	-42	7	-35	32	18	97	2	4	6

green = Δ > 50; yellow = Δ ≥ 30 & Δ < 50; orange = Δ > 10 & Δ < 30; gray = Δ < 10

Findings and Ongoing Work

Sensitivity Analysis Findings

- *Koala* cloud simulator exhibits 8 behavioral dimensions
- 6 input parameters significantly influence *Koala* behaviors
- Using a 2-level experiment design, comparison of resource allocation heuristics will require no more than (2^6 =) 64 parameter combinations, fewer (e.g., 2^{5-1} = 32) with 2nd application of 2-level OFF experiment design

- Analysis of direction in response correlations and of direction and magnitude in parameter effects suggests *Koala* behaviors are sensible
- *Koala* resource requirements permit simulation of moderately-sized cloud configurations O(10^5) nodes
- *Koala* implementation and our computing infrastructure appear robust enough for maximum simulation durations lasting months

Ongoing Work

Currently conducting an experiment to compare 18 resource allocation heuristics for on-demand IaaS Clouds

Cluster Selection	Node Selection
Least Full First	First Fit
	Next Fit
Percent Allocated	Tag & Pack
	Random
Random	Least Full First
	Most Full First
3 x	6 = 18

Experiment design is "Resolution VI" 2^{5-1} OFF, requiring simulating each of the 18 heuristics under 32 conditions (i.e., 576 total simulations)

Simulations are completed, data collected and analyzed. Paper in preparation.

Backup Slides

Koala Information Visualizations by Sandy Ressler

(see http://math.nist.gov/~SRessler/cloudviz.html for animations and more)

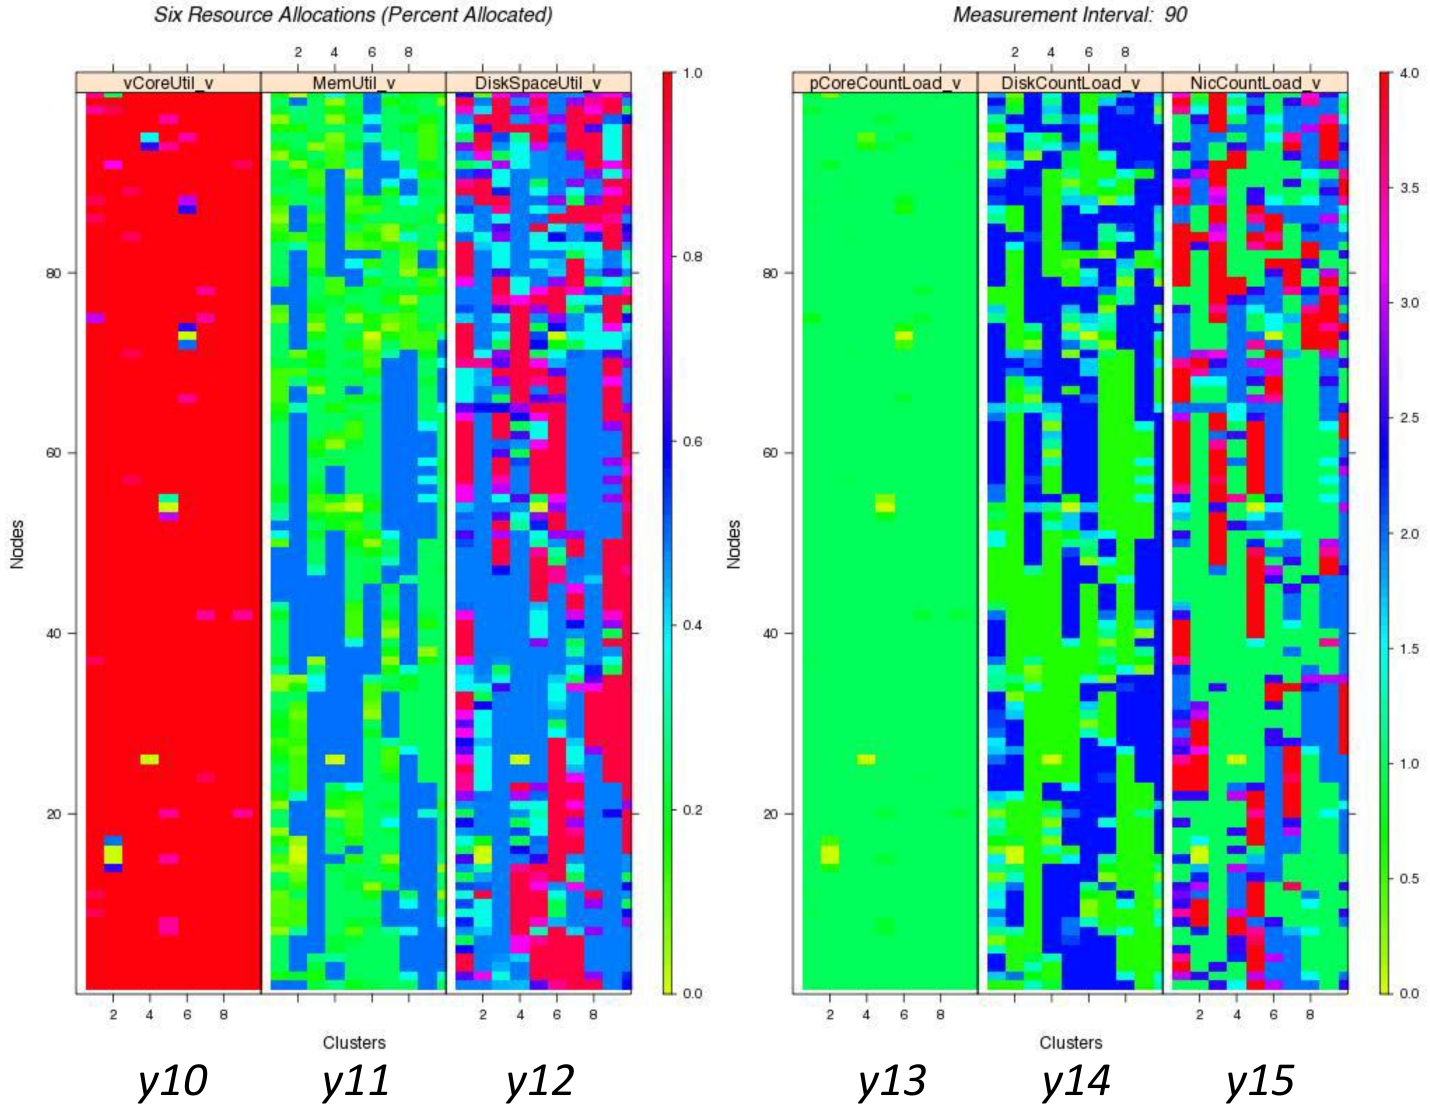

Koala Performance Characteristics for Sensitivity Analysis Experiments

Resource Usage: Time & Space

Experiment	Time (minutes) μ	σ	min	max	Memory (Megabytes) μ	σ	min	max
SA1-small (avg. over 6 reps.)	17	13	2	69.3	54	12	37	73
SA2-small (avg. over 6 reps.)	56	148	<1	1019	70	44	27	187
SA1-large	2389	2723	211	12,645	266	91	134	467
SA2-large	4173	7659	29	38,057	235	179	53	764

Simulated Configurations: Time & Space

Experiment	Simulated Time (Hours) MINUS level	PLUS level	Number of Users MINUS level	PLUS level	Number of Nodes MINUS level	PLUS level
SA1-small	600	1200	250	500	1000	4000
SA2-small	200	1600	125	750	250	12,000
SA1-large	600	1200	2500	5000	10,000	40,000
SA2-large	200	1600	1250	7500	2500	60,000