

OCCAM: Is there a future for open-access simulation and experimentation?

Bruce Childers†, Alex Jones‡, **Daniel Mossé†**


†Computer Science Department

‡Electrical & Computer Engineering

University of Pittsburgh, Pittsburgh, PA 15260

{childers,mosse}@cs.pitt.edu


akjones@ece.pitt.edu



Supported by



Computer Architecture Innovation




- \$111B processor market[†]
- Relies almost exclusively on artifacts
 - Software simulation
 - Hardware emulation
 - Benchmarks
 - And a cast of a thousand other tools...




[†] IMS Research, 2011

Prolific Artifact Production



	Simulators																															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
Single-core	X	X	X	.	.	X	X	X	X	X	X	X	X	X	X	X	X	.	X	X	X	.	X	X	X	X	X	X	X	X	X	X
Multi-core	X	X	X	.	.	X	X	X	X	X	X	.	.	X	X	.	.	.	X	X	X	X	.	X	X	.	.	X	.	.	.	
Homogen. multi-core		X	X	.	.	X	X	X	X	X	X	.	.	X	X	.	.	.	X	.	X	X	X	X	.	.	.	X	.	.	.	
Heterogen. multi-core	X	X	
SMT				.	.			X		X				X	X	X		X	
Shared memory		X	X	.	.	X	X	X	X	X		.	.	X	X	X	X	.	X	X	X	X	X	X	X	
Private memory	
Timing	X		X	X	.	X	X	X	X	X	X	X	X	X	X	X	.	X	X	.	X	X	X	X	.	X	X	.	X	X	.	
Cycle-accurate	.		X	X	.	X	X	X	X		X	X	X	X	X	X	.	X	.	X	X	X		.	X	X	.	X	X	X	X	
Functional	.	X	.	.	.		X	X	X	X	X	X	X	X	X	.	.	.	X	.	X	.	.	X	X	.	X	.	X	.	.	
Full system	X	X	.	.	X		X	X	X	X	.	.	X	X	X	X	
Caching	.	.	X	X	.	X	X	X	X	X	X	X	X	X	X	X	.	X	.	X	X	X	X	X	X	.	X	X	X	.	X	
Cache coherence	.	.			.	X	X	X	X	X	.	.	X	X		X	.	X	.	X	.	X	X	.	.	.	X	
In-order	X		X	X	X	X	X	X	X	X	.	X	.	X	.	X	X	.	.	.	X	X	.	X	X	X	
Out-of-order	X	X	X	X		X	X	X	X	X	.	.	X	.	.	.	X	X	.	.	X	.	X	.	.	.	
Superscalar	X	X	X	X	X	X	.	X	X	X	.	.	X	.	.	X	X	X	.	.	X	.	X	.	.	.	
Virtualization	X	
Virtual memory	.	X	.	.	X		X		X		.	X	X	X		.	X	X	X	X	
VLIW	X	X	
DRAM controller	X	X	X	X	X	.	X	X	.	X	.	.	.	X	X	X	.	.	.	X	.	X	
Scheduling		X	X	X	X	.	X	X	X	X	X	.	.	.	X	.	X	
Concurrency		X	X	X	.	.	X	X	X	X	X	
DRAM error sim	
On-chip network	X	X	
System on chip	X	
Power consumption	.	.	X	X	.	.	.	X		X	.	X	.	.	.	X	X	X	.	.	X	.	.	
Gate-level	X	X	.	.	
Validated	X		X	.	X	X	.	X	.	.	.	X	X	X	.	.	




Prolific Artifact Production

	Simulators																															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
Single-core	X	X	X	.	.	X	X	X	X	X	X	X	X	X	X	X	X	.	X	X	X	.	X	X	X	X	X	X	X	X	X	X
Multi-core	X	X	X	.	.	X	X	X	X	X	X	.	.	X	X	.	.	.	X	X	X	X	.	X	X	.	.	X	.	.	.	
Homogen. multi-core		X	X	.	.	X	X	X	X	X	X	.	.	X	X	.	.	.	X	.	X	X	X	X	.	.	.	X	.	.	.	
Heterogen. multi-core	X	X	
SMT				.	.			X		X				X	X	X			X	
Shared memory	X	X	.	.	X	X	X	X	X		.	.	X	X	X	X	.	X	X	X	X	X	X	X	X	
Private memory	
Timing	X		X	X	.	X	X	X	X	X	X	X	X	X	X	.	X	X	.	X	X	X	X	X	.	X	X	.	X	X	.	
Cycle-accurate	.		X	X	.	X	X	X	X		X	X	X	X	X	X	.	X	.	X	X	X		.	X	X	.	X	X	X	X	
Functional	.	X	.	.	.		X	X	X	X	X	X	X	X	.	.	.	X	.	X	.	.	X	X	.	X	.	X	.	.	.	
Full system	X	X	.	.	X		X	X	X	X	.	.	X	X	X	X	
Caching	.	.	X	X	.	X	X	X	X	X	X	X	X	X	X	X	.	X	.	X	X	X	X	X	X	.	X	X	X	.	X	
Cache coherence	.	.			.	X	X	X	X	X	.	.	X	X		X	.	X	.	X	.	X	X	.	.	.	X	
In-order							X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	
Out-of-order																																
Superscalar																																
Virtualization																																
Virtual memory																																
VLIW																																
DRAM controller																																
Scheduling																																
Concurrency																																
DRAM error sim																																
On-chip network																																
System on chip																																
Power consumption	.	.	X	X	.	.	.	X				X	.	.	.	X	X	X	.	.	X	.	.	
Gate-level		X	X	.	.	
Validated	X		X	.	X	X	.	X	.	.	X	X	X	.	.	


Just a *small* selection of 31 artifacts!
 Tremendously diverse, different, overlapping...
 Compounded by many experiments...

A brewing (brewed?) crisis...




- Symptomatic of fragmented, ad hoc, internal existing effort and investment
- Research expediency and results
 - Small incentive to build, release, maintain
 - Duplicating effort to re-implement for comparison
 - May not understand artifact and its use
 - Creating your own artifacts instead
- Missing (often?) accountability, repeatability

Community, open source effort




- Community **is asking for openness**
 - Activities by NSF, DOE, DARPA/DoD, CRA over last ten years to inspire, engage the community
 - More recent activities events at SC11—SC13, MICRO45, HiPEAC CSE13, ISCA13, and today!
- Quality artifacts **are** available, emerging
- Quality experiments **are** being done
- How can we build and leverage this?



Vision for OCCAM

- Inspired by Occam's Razor, which suggests minimum assumptions and most succinctness
- Create a shared instrument that everyone can and should use to save time, be fair, advance science: *OCCAM: Open Curation for Computer Architecture Modeling*
- Community-supported digital curator for simulation, emulation benchmarking and experimental results






Vision for OCCAM

- Efforts in different areas successful
 - nanoHUB (Nanotechnology)
 - arXiv (Physics and other sciences)
 - EarthCube (Geosciences)
 - C-tuning
 - Data-mill
 - Others

OCCAM pillars

Building a bridge to open access

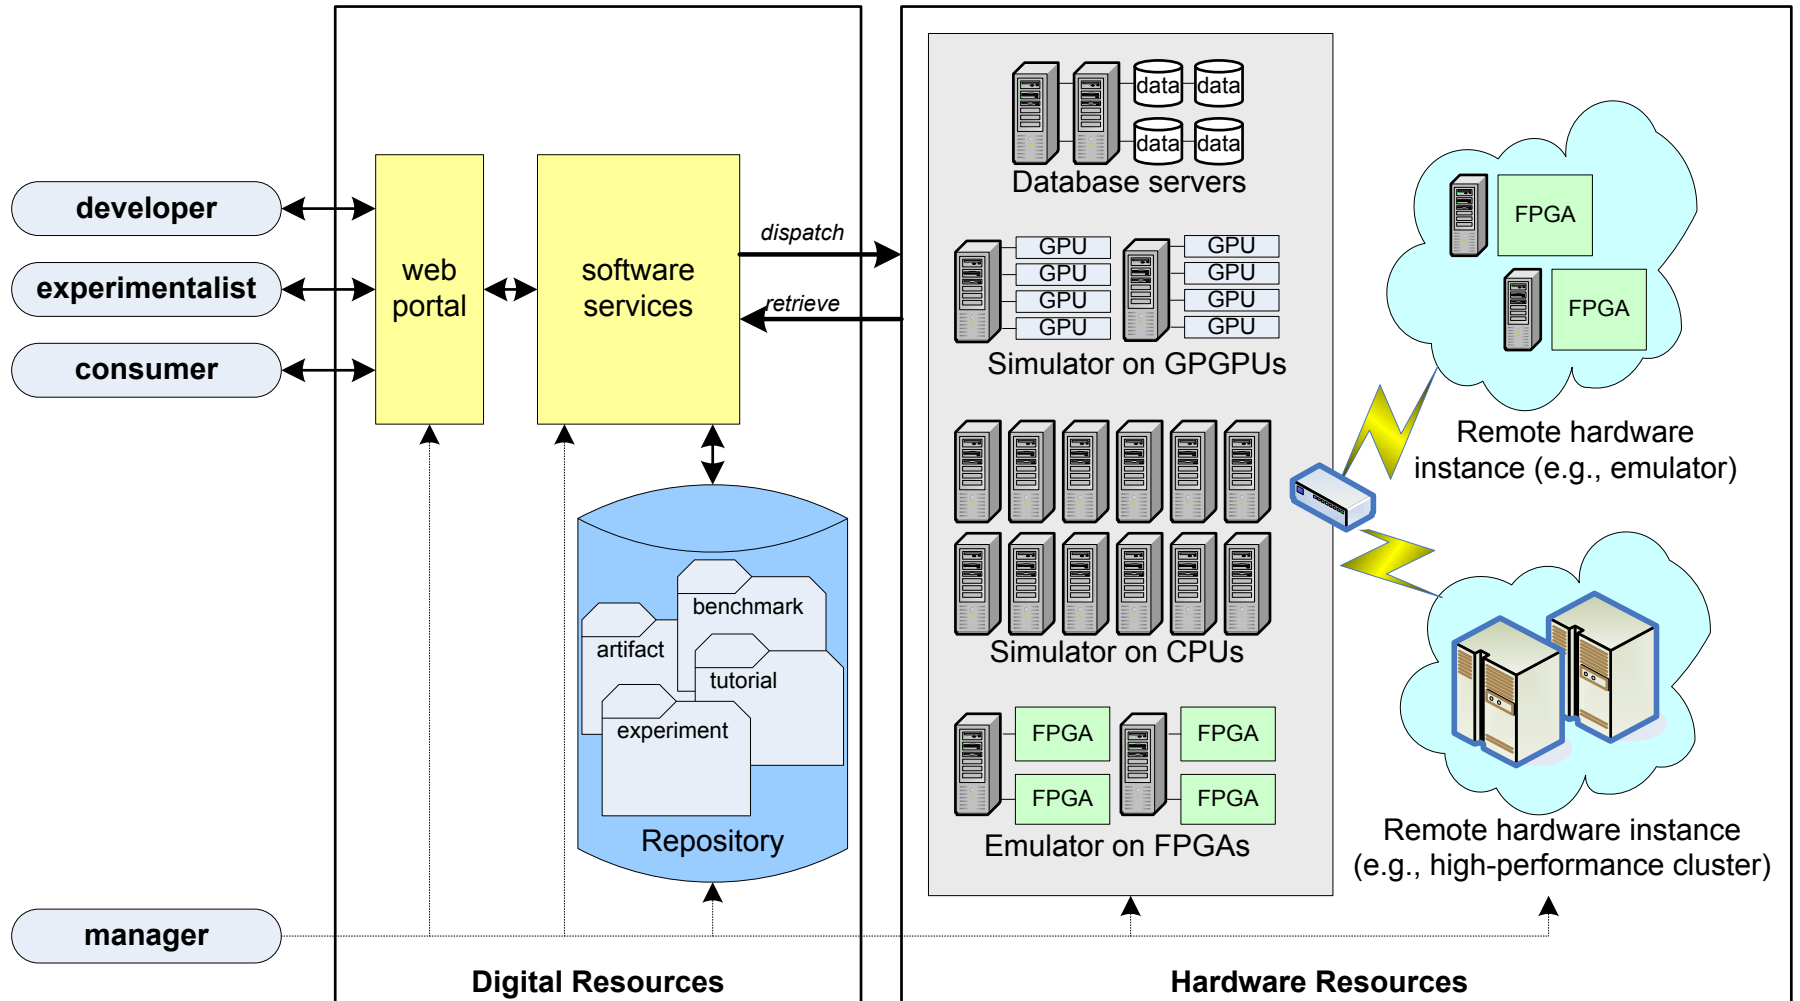



OCCAM pillars

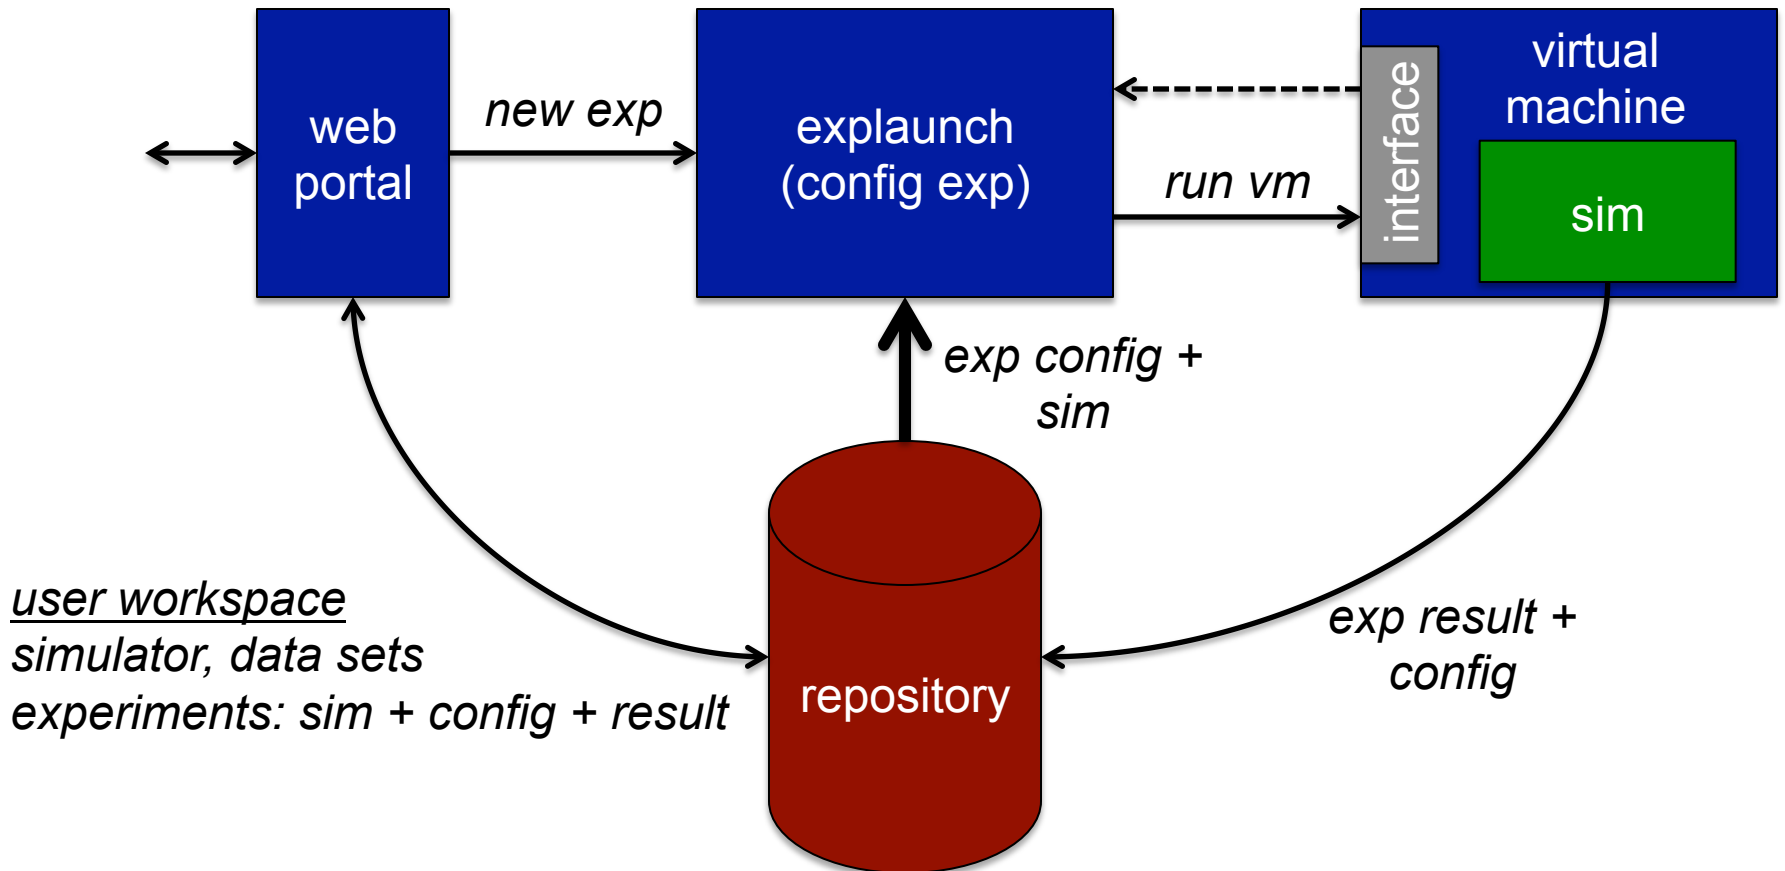

Building a bridge to open access



Infrastructure Pillar



Infrastructure Pilot (“v0”)







Community Pillar

- Create, foster community
 - Establish procedures and policies
 - Contribute artifacts and experiments
 - Ingrain and use repository, e.g., conference submission and publication of experiments
 - Advocate and convince people of value and values behind open-access repository

Community Pilot (“v0”)



- Determining governance structures
- Initial procedures and policies
 - ① Incentive and reward
 - ② Privacy and security (privacy, data/code safety, access rights, time limits)
 - ③ Resource contribution (credit/pricing, fairness)
 - ④ Methodology (guidelines, checks, rewards)
 - ⑤ Assessment (usage, quality of service to users)
 - ⑥ Ontology (classification for query & search)



Education Pillar

- “Build it and they’ll come” is not a paradigm that has worked
- Rather, address *barrier to adoption*
 - Best practices of building artifacts, using them and running experiments
 - Scientific process (accountable, repeatable)
- Training
 - Developer’s, Experimentalist's Bootcamp
 - Summer School for Simulation and Emulation
 - Short (10 min), Medium (30-60 min), Long (hrs)


Starting the Effort



- Pilot repository <http://www.occamportal.org>

News! Content!

1. **Survey of needs**
2. Best practices
4. Catalog of artifacts



It Takes a Community...

Many++ have contributed... CSA workshop and other activities

Tor Aamodt
Ahmed Amer
Jason Bakos
Christopher Batten
Nathan Binkert
Bruce Childers
Derek Chiou
Sangyeun Cho
Almadena Chtchelkanova
John Davis
Sandhya Dwarkadas
Lieven Eeckhout
Kriszitian Flautner
Jean-Luc Gaudiot
James Hoe

Engin Ipek
Alex Jones
Hyesoon Kim
Martha Kim
Michael Kistler
Jack Lange
Benjamin Lee
Ahmed Louri
Nicolas Maillard
Jason Mars
Chris Mineo
Daniel Mossé
Frank Mueller
Steve Poole
Steven Reinhardt
Arun Rodrigues

Eric Rotenberg
Sonia Sachs
Ali Saidi
Kevin Skadron
Evan Speight
Richard Uhling
James Tringali
Noel Wheeler
David Wood
Sudhakar Yalamanchili
Li Zhao
SC11 BoF
SC12 BoF
MICRO45
HiPEAC CSE13
ISCA13
MSE

Thank you