


Repeatable research with open source platforms



Andrew W. Moore, PhD CEng
Computer Laboratory
University of Cambridge

<http://www.cl.cam.ac.uk/~awm22/slides/2017-coseners.pptx>



UNIVERSITY OF
CAMBRIDGE

Reproducibility in Science

- **Validate Correct Results**

supporting the conclusions
and
compare with new ideas

- **Invalidate Incorrect results**

refuting the conclusions
and
improve and refine

Reproducibility as validation



Reproducibility as invalidation



The image is a screenshot of a news article on the Nature website. At the top, the word "nature" is written in a large, white, serif font on a dark red background, with the subtitle "International weekly journal of science" in a smaller, white, sans-serif font below it. A navigation bar below the header contains several white buttons with red text: "nature news home", "news archive", "specials", "opinion", "features", "news blog", and "na".

On the left side, there is a blue speech bubble icon with an exclamation mark, followed by the text "comments on this story" in red and blue. Below this is a dashed line and the text "Stories by subject". Underneath, there is a red bullet point followed by the word "Physics" in red.

The main content area on the right starts with the text "Published online 22 September 2011 | Nature | doi:10.1038/news.2011.554" and "Updated online: 23 September 2011". Below this is a red box with the word "News" in white. The main headline is "Particles break light-speed limit" in a large, bold, black serif font. Below the headline is a dashed line and the sub-headline "Neutrino results challenge cornerstone of modern physics." in a bold, black sans-serif font.

Reproducibility as invalidation

The image is a screenshot of the Nature website. At the top left, the word "nature" is written in a large, white, serif font, with the subtitle "International weekly journal of science" in a smaller, white, sans-serif font below it. Below this, there is a navigation bar with several tabs, including "nature news". A "comments on story" link is visible. The main content area features the "Science" logo in a large, white, serif font, with the AAAS logo to its right. Below the logo is a red navigation bar with the following tabs: "Home", "News", "Journals", "Topics", and "Careers". Underneath this bar is a secondary navigation bar with the following links: "Latest News", "ScienceInsider", "ScienceShots", "Sifter", "From the Magazine", "About News", and "Quizzes".

SHARE

Once Again, Physicists Debunk Faster-Than-Light Neutrinos

By **Adrian Cho** | Jun. 8, 2012 , 3:39 PM

What do Linux apache MySQL Firefox BSD BIND



have in common with the resurgence in
Software Defined Networking (SDN)?

What do Linux apache MySQL Firefox BSD BIND



have in common with the resurgence in SDN?


Openly available source and open standards



And not just an implementation...

- OFLOPS – the OpenFlow performance tester
- OFtest – OpenFlow compliance tester

- Ovs – software only implementation



- Numerous OpenFlow controllers:

From Ryu to OpenDaylight





Each a stable platform

- enabling extension, and
- a process for adopting contributions and improvements




Because **network research**, and **education**
needs a good **platform**

www.netfpga.org

So what is NetFPGA?


NetFPGA = Networked FPGA platform


A line-rate, flexible, open networking platform for teaching and research

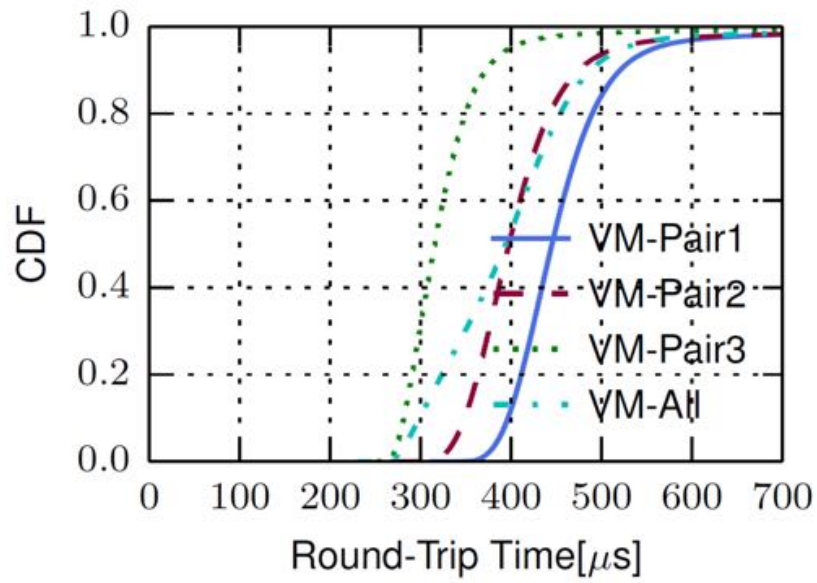


MIPS, RISC-V, Blueswitch, EMU, P4 FPGA...


Traffic Control – The Next Gen

- Latency injection – ns to seconds granularity
- Different latency distributions
- Rate control
 - Rate limiting
 - Shaping
- Per flow support
- Powered by  **NetFPGA**







Datacentre Latency




Datacentre Latency



Small latencies can still lead to significant performance loss




Measure the Application


- Packet Size
- Inter-Packet Gap
- Burst Size
 - Is congestion likely?
- Bandwidth – at different granularities
- Powered by  NetFPGA

Bandwidth – 10ms granularity


Lasso (ML)



Memcached




Apache




Bandwidth – 100us granularity


Lasso (ML)



Memcached

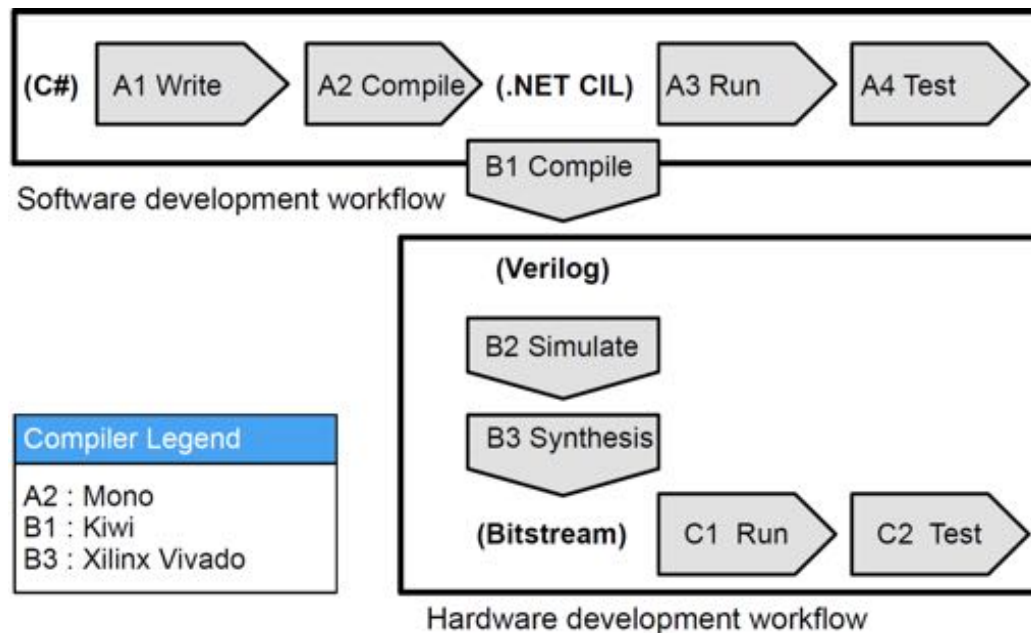


Apache



Emu : Accelerating Network Services

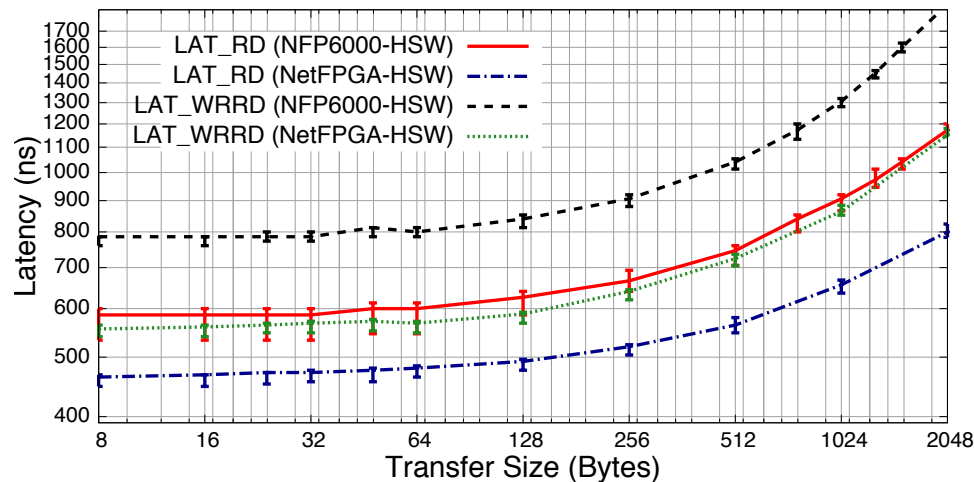
- The network library for Kiwi
- Compiling .Net programs
 - To x86
 - To simulation environment
 - To multiple FPGA targets



pcie-bench: an open source tool for benchmarking PCI Express

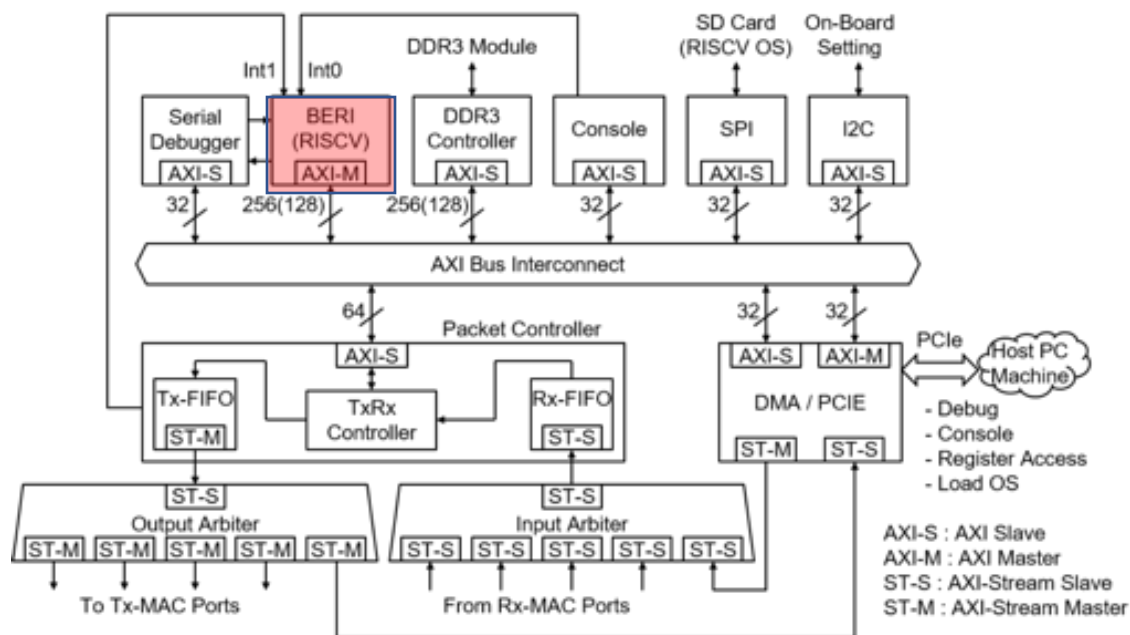
there is a limited understanding of PCIe functionality, nor the trade-offs that must be made to get best-performance from PCIe systems

- pcie-bench tool open source available
- It builds on NetFPGA and Netronome boards



NetSoc: NetFPGA + Open Source Processors

- Open source, RISC based SoC architectures
- With high performance networking capabilities
- *RISC-V* – *RISC-V ISA soft processor, Linux OS*
- *CHERI* – *64bit MIPS soft processor, BSD OS*



What can we do?

- Raise the bar on acceptable research
 - Insist on the artefacts being published
 - Insist on the results being available
 - Insist on the experiments being repeatable
 - Accept reproduction studies
- Enable Research Repeatability

Lead by example

Lead by example

Moore & Zuev 2005, we published the dataset

```
https://www.cl.cam.ac.uk/research/srg/netos/projects/archive/nprobe/data/papers/sigmatrics/index.html
```

In a form that was preserved anonymity and

1. Enabled reproducibility,
2. Enabled comparison against new algorithms, and
3. Actively encouraged others to create their own datasets

Lead by example

Queues don't matter when you can JUMP them!

Matthew P. Grosvenor Malte Schwarzkopf Ionel Gog Robert N. M. Watson
Andrew W. Moore Steven Hand Jon Crowcroft

University of Cambridge Computer Laboratory

Simplicity is the shortest path to a solution.
– Ward Cunningham

Abstract

QJUMP is a simple and immediately deployable approach to controlling network interference in datacenter networks. Network interference occurs when congestion from throughput-intensive applications causes queueing that delays traffic from latency-sensitive applications. To mitigate network interference, QJUMP applies Internet

If memcached packets can somehow be prioritized to “jump-the-queue” over Hadoop’s packets, memcached will no longer experience latency tails due to Hadoop. Of course, multiple instances of memcached may still interfere with *each other*, causing long queues or incast collapse [10]. However, if each memcached instance can be appropriately rate-limited at the origin, this too can be mitigated.

These observations are not new: QoS technologies like DiffServ [7] demonstrated that coarse-grained classification and rate limiting can be used to control network

Lead by example

Queues don't matter when you can JUMP them!

<http://www.camsas.org/qjump>



Lead by example

Queues don't matter when you can JUMP them!

<http://www.camsas.org/qjump>

QJump

Learn

Publications

Download

Reproduce

About Us



CamSaS

Reproducing the QJump Experiments








As scientists and researchers we take the reproducibility of our work very seriously. We don't expect you to trust our results, in fact, we hope that you don't! On the following page we provide links to detailed descriptions the experiments behind each of the figures in our research publications. The full experimental descriptions include the precise configuration of our test equipment, links to the source code for our tools, patches that we made to other people's tools and original preprocessed data-sets that we gathered. Our hope is that anyone can use these descriptions re-run any of our experiments. We believe that this kind of openness is the way that all good, scholarly scientific research should be conducted.

For details of our publications including links to the manuscripts please see the [publications](#) page.

NSDI 2015 - Queues don't matter when you can












Enable others

Reproducible research needs,
widely available test-equipment

	Cost	Flexibility	Resolution	Line Rate
	\$\$\$ \$\$\$			
DPDK, SW tools	(\$)			()

Enable others

Reproducible research needs,
widely available test-equipment

	Cost	Flexibility	Resolution	Line Rate
	\$\$\$ \$\$\$			
DPDK, SW tools	(\$)			()
	(\$)			

Open-Source Network Tester



A platform for testing powered by



- Open source hardware and software platform for network test, publicly available

<https://osnt.org/>

<https://github.com/NetFPGA/OSNT-Public/wiki>

- Low cost, low jitter, flexible to update, scale-out, no CPU usage, nano-second resolution measurement

What can we do?

- Raise the bar on acceptable research
 - Insist on the artefacts being published
 - Insist on the results being available
 - Insist on the experiments being repeatable
 - Accept reproduction studies
- Enable Research Repeatability
and then *Just Do It!* Everytime.
- Open-source (research) platforms ***work!***
Build on platforms, and support platforms

What is the platform for machine-learning in networking?

Is the lack of an *active*-optical network due to a lack of platform?
(no one can make up their mind what it looks like...?)

What is the next software-platform to displace the current one?

(So what is Ion Stoica's next project 😊...?)

Acknowledgements



UNIVERSITY OF
CAMBRIDGE

EPSRC

Pioneering research
and skills



XILINX

ALL PROGRAMMABLE™



The Leverhulme Trust



HUAWEI

ALGO-LOGIC

