Research Paper Reading Group

Pilot Session 2

Previous session

Types

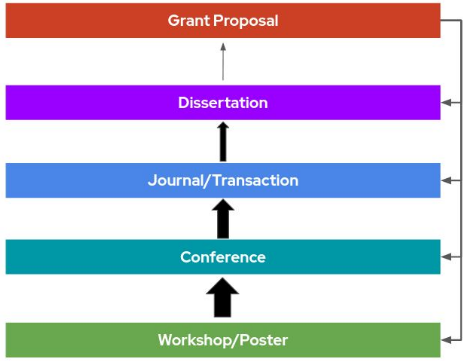

Tiers

	Avg citations per paper	Conference
1.	66.3	CSUR—ACM Computing Surveys
2.	53.5	SOSP—ACM Symposium on Operating Systems Principles
3.	52.8	OSDI—Operating Systems Design and Implementation
4.	43.4	NDSS—Network and Distributed System Security Symposium
5.	40.9	MobiHoc—Mobile Ad Hoc Networking and Computing
6.	40.3	SIGCOMM—ACM SIGCOMM Conference
7.	38.2	SenSys—Conference On Embedded Networked Sensor Systems
8.	36.7	MOBICOM—Mobile Computing and Networking
9.	36.1	CIDR—Conference on Innovative Data Systems Research
10.	35.3	USENIX Security Symposium
11.	35.3	EUROCRYPT—Theory and Application of Cryptographic Techniques
12.	35.0	NSDI—Networked Systems Design and Implementation
13.	34.4	JASSS—The Journal of Artificial Societies and Social Simulation
14.	33.5	TOCS—ACM Transactions on Computer Systems
15.	33.5	S&P—IEEE Symposium on Security and Privacy
16.	33.4	MobiSys—International Conference on Mobile Systems
17.	32.5	IJCV—International Journal of Computer Vision
18.	32.2	TOG—ACM Transactions on Graphics/SIGGRAPH
19.	31.6	VLDB—Very Large Data Bases
20.	30.9	BioMED—Biomedical Engineering
21.	30.9	IEEE TRANS ROBOTICS AUTOMAT—IEEE Transactions on Robotics and Automation
22.	30.6	CRYPTO—International Crytology Conference
23.	30.1	PAMI—IEEE Transactions on Pattern Analysis and Machine Intelligence
24.	29.6	PLDI—SIGPLAN Conference on Programming Language Design and Implementation
25.	29.3	MICRO—International Symposium on Microarchitecture
26.	29.1	Journal of Web Semantics
27.	28.5	BIB—Briefings in Bioinformatics
28.	27.4	JMLR—Journal of Machine Learning Research
29.	27.0	ISMB—Intelligent Systems in Molecular Biology
30.	26.8	PODS—Symposium on Principles of Database Systems
31.	26.5	VTC—Vehicular Technology Conference
32.	25.5	SIGMOD—International Conference on Management of Data
33.	24.6	STOC—ACM Symposium on Theory of Computing
34.	24.0	TOIS—ACM Transactions on Information Systems
35.	24.0	IEEE SAP—IEEE Transactions on Speech and Audio Processing
36.	23.9	SCA—Symposium on Computer Animation
37.	23.3	BIOINFORMATICS—Bioinformatics/computer Applications in The Biosciences

Format

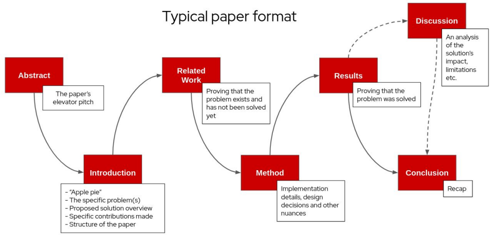

Value

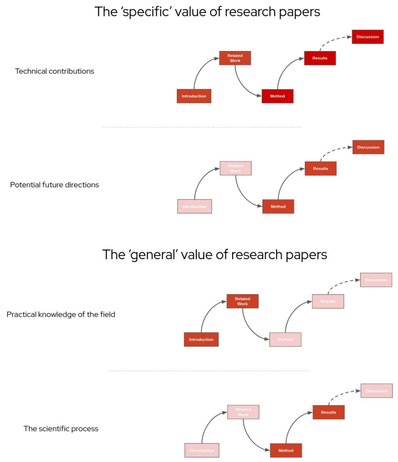

Pilot overview

Session 1	4/3/2021	A perspective on research papers
Session 2	5/4/2021	Identifying worthwhile papers
Session 3	6/1/2021	Discussing research papers

Common indicators

Citations	Higher citation count typically indicates higher quality work.

Red Hat

Citations

Deep Residual Learning for Image Recognition

Kaiming He Xiangyu Zhang Shaoqing Ren Jian Sun
Microsoft Research
{kahe, v-xiangz, v-shren, jiansun}@microsoft.com

Total citations

Cited by 76533

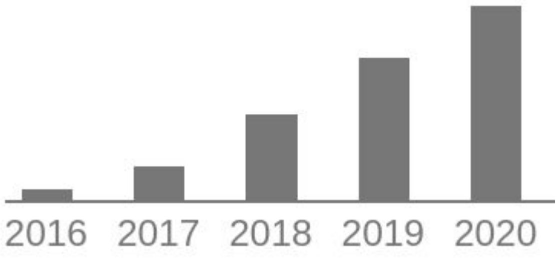

PROC. OF THE IEEE, NOVEMBER 1998 1

Gradient-Based Learning Applied to Document Recognition

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner

Total citations

Cited by 35879

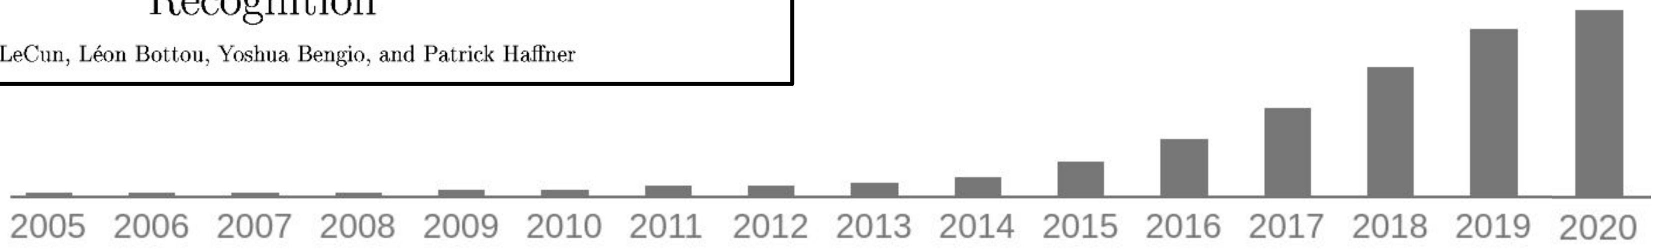

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.

Authors

Cited by

	All	Since 2016
Citations	198044	191726
h-index	57	56
i10-index	65	65

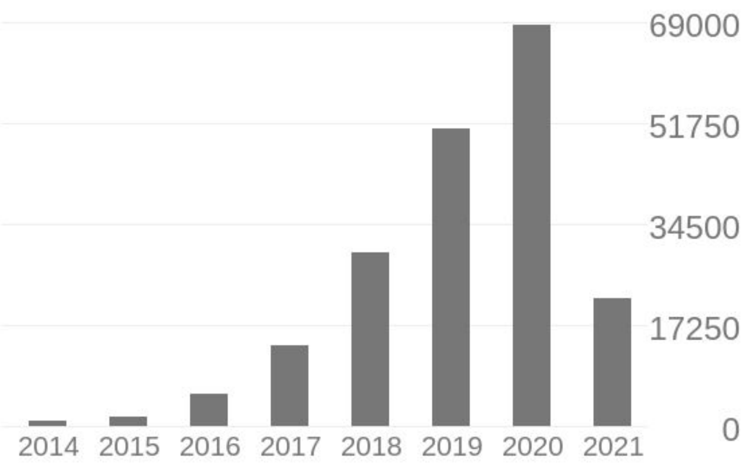

Cited by

	All	Since 2016
Citations		
h-index		
i10-index		

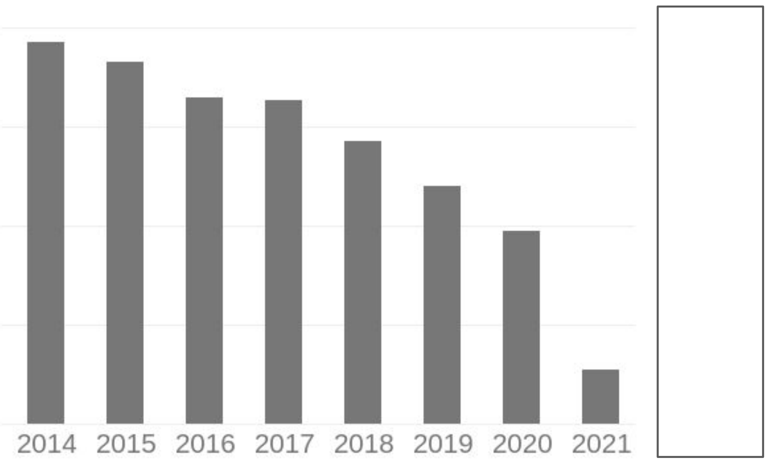

If the paper has too many authors, focus on the first and last one.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.

Scientific process

Providing SLOs for Resource-Harvesting VMs in Cloud Platforms

Characterizing unallocated resources: Azure 2/19 - 10/19 data	Methodology, Temporal patterns, Cluster behaviors, Regional aggregated data, Minimum unallocated cores, Additional unallocated cores, Multiple VMs per server, High-level takeaways
Proposed VM class: Harvest VM	Overview, Production implementation approach, Comparison to standard evictable VMs, Workload/application requirements, Privacy/Confidentiality, Pricing, Harvesting resources other than cores
Prediction for survival rate: ML based approach	User input, ML models and features, ML training and inference, Discarded features, Applying prediction to standard evictable VM survivability
Scheduler support: Harvest Hadoop	Architecture, Eviction management, Core reassignment management, Harvesting resources other than cores
Evaluation	Evaluation focus, Simulator, Experiments, Analysis (benefits, accuracy, scheduler, cost)

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.
System diagrams	Most good papers have one or more block diagrams / flow charts that illustrate their work and its different pieces.

System diagrams

Providing SLOs for Resource-Harvesting VMs in Cloud Platforms

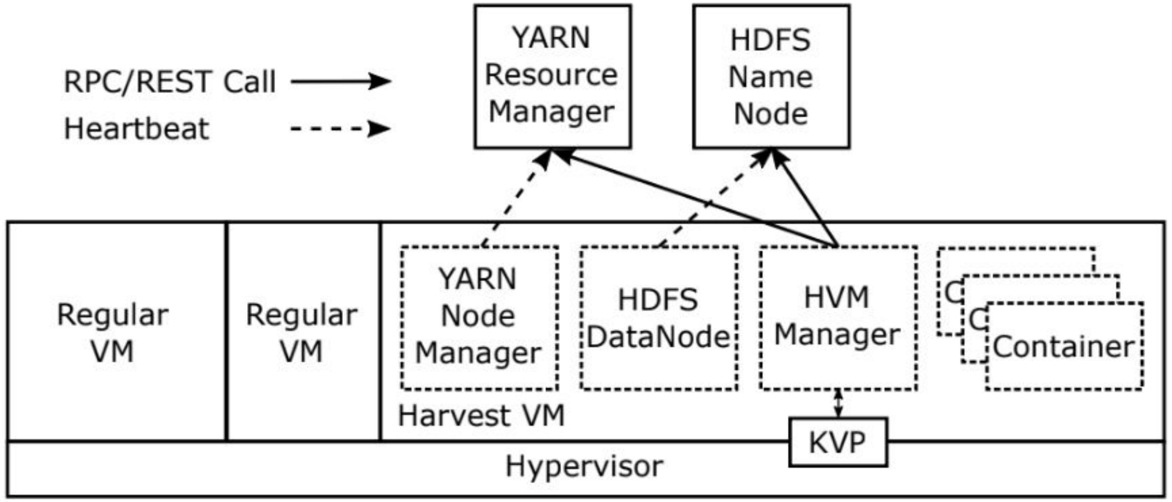

Red Hat

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.
System diagrams	Most good papers have one or more block diagrams / flow charts that illustrate their work and its different pieces.
Total figures	More illustrations can correspond to more work done (unless it is a survey paper).

Total figures

PANIC: A High-Performance Programmable NIC for Multi-tenant Networks

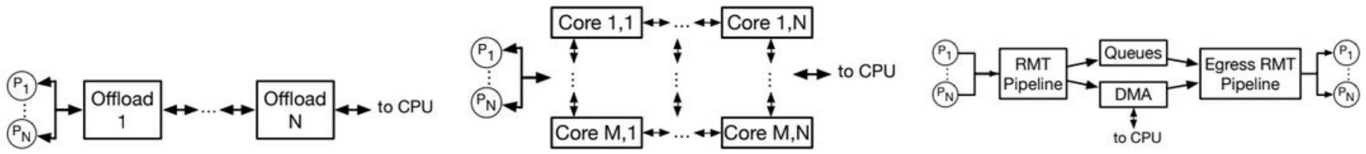

(a) **A pipeline-of-offloads architecture** (b) **A tiled manycore NIC architecture** (c) **A NIC with RMT pipeline [42]**

Figure 1: **Illustrations of existing programmable NIC architectures.**

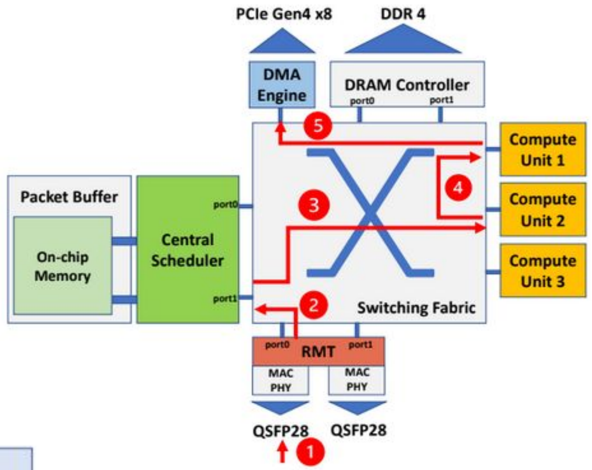

Figure 2: **PANIC Architecture**

Figure 3: **PANIC Descriptor**

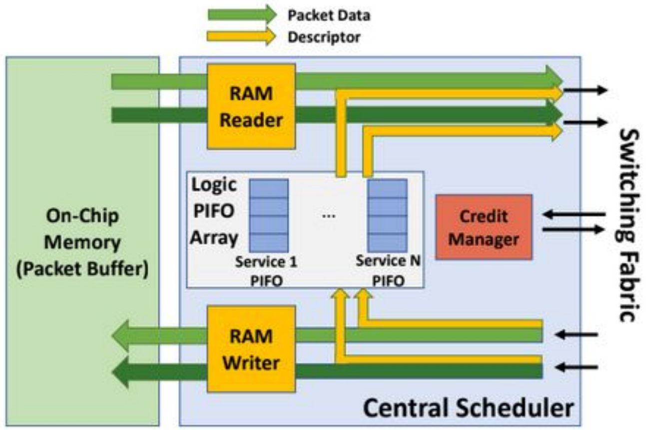

Figure 4: **Architecture of the multi-ported central scheduler**

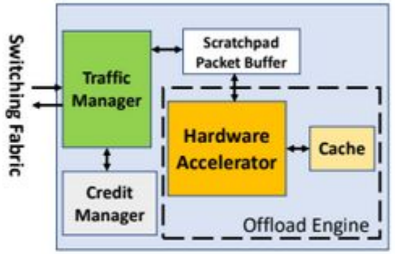

Figure 5: **Accelerator-Based Compute Unit Design**

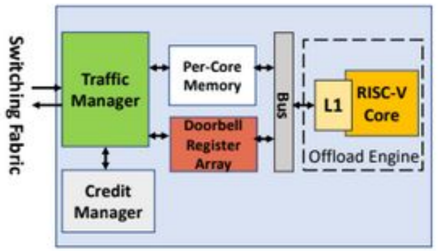

Figure 6: **Core-Based Compute Unit Design**

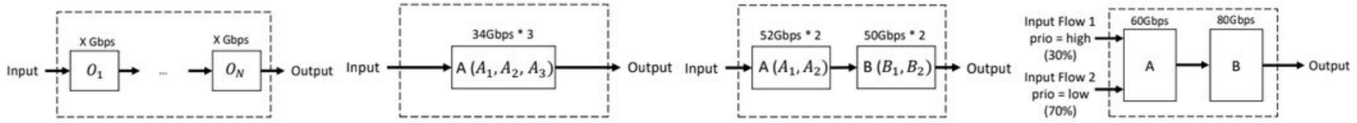

(a) **Chaining Model 1** (b) **Chaining Model 2** (c) **Chaining Model 3** (d) **Chaining Model 4**

Figure 7: **The different chaining models used in experiments.**

Red Hat

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.
System diagrams	Most good papers have one or more block diagrams / flow charts that illustrate their work and its different pieces.
Total figures	More illustrations can correspond to more work done (unless it is a survey paper).
Experiments done	Enough experiments should be done to demonstrate that the proposed contribution has been made.

Experiments done

A Framework for Acceleration of CNN Training on Deeply-Pipelined FPGA Clusters with Work and Weight Load Balancing

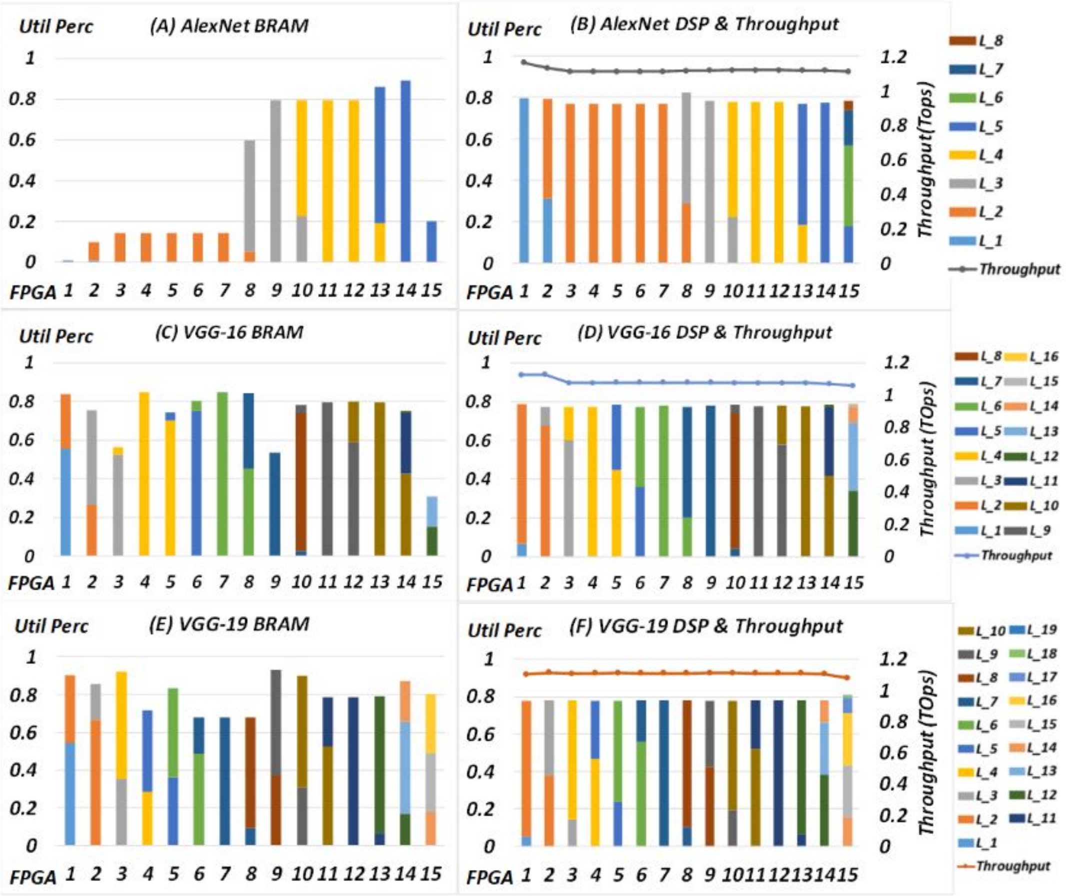

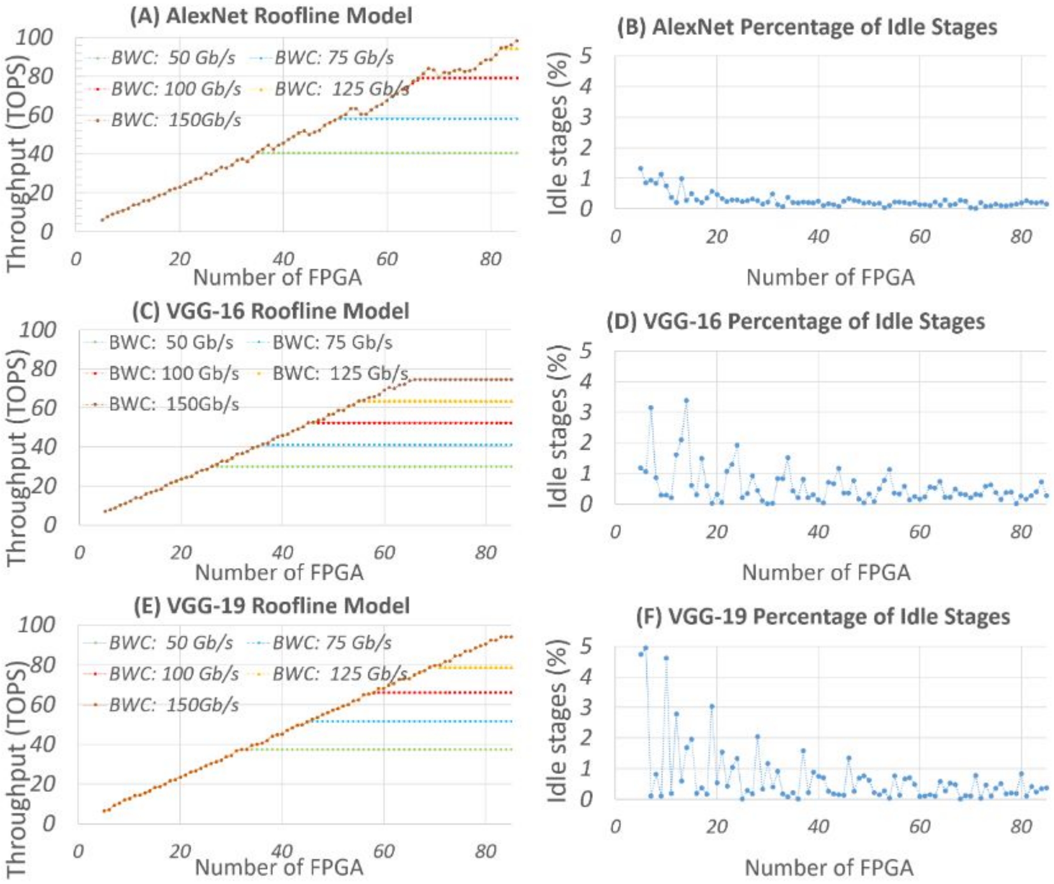

Red Hat

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.
System diagrams	Most good papers have one or more block diagrams / flow charts that illustrate their work and its different pieces.
Total figures	More illustrations can correspond to more work done (unless it is a survey paper).
Experiments done	Enough experiments should be done to demonstrate that the proposed contribution has been made.
Test environment	There can be a massive difference between testing something in theory, in simulation and in implementation on real hardware.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.
System diagrams	Most good papers have one or more block diagrams / flow charts that illustrate their work and its different pieces.
Total figures	More illustrations can correspond to more work done (unless it is a survey paper).
Experiments done	Enough experiments should be done to demonstrate that the proposed contribution has been made.
Test environment	There can be a massive difference between testing something in theory, in simulation and in implementation on real hardware.
Formatting	Good papers typically go through many rounds of internal review before final submission and thus tend to have better formatting.

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.
System diagrams	Most good papers have one or more block diagrams / flow charts that illustrate their work and its different pieces.
Total figures	More illustrations can correspond to more work done (unless it is a survey paper).
Experiments done	Enough experiments should be done to demonstrate that the proposed contribution has been made.
Test environment	There can be a massive difference between testing something in theory, in simulation and in implementation on real hardware.
Formatting	Good papers typically go through many rounds of internal review before final submission and thus tend to have better formatting.
Open source	If the code is open source and reproducible, the credibility of the work increases substantially

Case study

Paper 1

A Reconfigurable Fabric for Accelerating Large-Scale Datacenter Services

Andrew Putnam Adrian M. Caulfield Eric S. Chung Derek Chiou[1]
Kypros Constantinides[2] John Demme[3] Hadi Esmaeilzadeh[4] Jeremy Fowers
Gopi Prashanth Gopal Jan Gray Michael Haselman Scott Hauck[5] Stephen Heil
Amir Hormati[6] Joo-Young Kim Sitaram Lanka James Larus[7] Eric Peterson
Simon Pope Aaron Smith Jason Thong Phillip Yi Xiao Doug Burger

Microsoft

https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/Catapult_ISCA_2014.pdf

Paper 2

A Cloud-Scale Acceleration Architecture

Adrian M. Caulfield Eric S. Chung Andrew Putnam
Hari Angepat Jeremy Fowers Michael Haselman Stephen Heil Matt Humphrey
Puneet Kaur Joo-Young Kim Daniel Lo Todd Massengill Kalin Ovtcharov
Michael Papamichael Lisa Woods Sitaram Lanka Derek Chiou Doug Burger

Microsoft Corporation

https://www.microsoft.com/en-us/research/wp-content/uploads/2016/10/Cloud-Scale-Acceleration-Architecture.pdf

Paper 3

The Feniks FPGA Operating System for Cloud Computing

Jiansong Zhang[§] Yongqiang Xiong[§] Ningyi Xu[§] Ran Shu[§†] Bojie Li[§‡]
Peng Cheng[§] Guo Chen[§] Thomas Moscibroda[§]
[§] Microsoft Research [†] Tsinghua University [‡] USTC
{jiazhang,yqx,ningyixu,v-ranshu,v-bojli,pengc,guoche,moscitho}@microsoft.com

https://www.microsoft.com/en-us/research/uploads/prod/2018/09/Feniks-APSys17.pdf

Scoring sheet

Indicator	Average Score	R1	R2	R3
Citations				
Year				
Authors/Organization				
Venue				
Abstract				
Specific contributions				
“Beyond the scope”				
Introduction length				
Scientific process				
System diagrams				
Total figures				
Experiments done				
Test environment				
Formatting				
Open source				

Paper 1

Indicator	Average Score	R1	R2	R3
Citations	1097			
Year	2014			
Authors/Organization	0	0	0	0
Venue	0	0	0	0
Abstract	10			
Specific contributions	10			
"Beyond the scope"	10			
Introduction length	0	0	0	0
Scientific process	10			
System diagrams	0	0	0	0
Total figures	0	0	0	0
Experiments done	0	0	0	0
Test environment	5			
Formatting	0	0	0	0
Open source	0			

Paper 2

Indicator	Average Score	R1	R2	R3
Citations	488			
Year	2016			
Authors/Organization	0	0	0	0
Venue	0	0	0	0
Abstract	10			
Specific contributions	10			
“Beyond the scope”	10			
Introduction length	0	0	0	0
Scientific process	10			
System diagrams	0	0	0	0
Total figures	0	0	0	0
Experiments done	0	0	0	0
Test environment	10			
Formatting	0	0	0	0
Open source	0			

Paper 3

Indicator	Average Score	R1	R2	R3
Citations	21			
Year	2017			
Authors/Organization	0	0	0	0
Venue	0	0	0	0
Abstract	10			
Specific contributions	10			
“Beyond the scope”	1			
Introduction length	0	0	0	0
Scientific process	10			
System diagrams	0	0	0	0
Total figures	0	0	0	0
Experiments done	0	0	0	0
Test environment	5			
Formatting	0	0	0	0
Open source	0			

Red Hat

Comparison

Indicator	P1	P2	P3
Citations	10	4.4	0.2
Year	10	10	10
Authors/Organization	0	0	0
Venue	0	0	0
Abstract	10	10	10
Specific contributions	10	10	10
“Beyond the scope”	10	10	1
Introduction length	0	0	0
Scientific process	10	10	10
System diagrams	0	0	0
Total figures	0	0	0
Experiments done	0	0	0
Test environment	5	10	5
Formatting	0	0	0
Open source	0	0	0
Overall score	65	64	46

Join us for the next session!

Session 1	4/3/2021	A perspective on research papers
Session 2	5/4/2021	Identifying worthwhile papers
Session 3	6/1/2021	Discussing research papers

Red Hat

Join us for the next session!

Session 3	6/1/2021	Discussing research papers

Homework: Read 1 research paper on "Deep learning for self driving cars" .

For every paper read, we will discuss:

1) The problem being targeted
2) The proposed solution
3) Benefits of the solution
4) Limitations of the solution

Responses can be shared in person during the session, or submitted beforehand via the QR code.

Red Hat

Sign-up / Comments / Suggestions / Feedback

https://forms.gle/6Y2ZBH2Bq2y5Qmie7

Thank you!