Computing without Servers, V8, Rocket Ships, and Other Batsh*t Crazy Ideas in Data Systems

Jimmy Lin
David R. Cheriton School of Computer Science
University of Waterloo

Wednesday, August 29, 2018

Source: Flickr (massimoleonardi/8064194094/)

DESIRES

A systems-oriented biennial conference, complementary in its mission to the mainstream Information Access and Retrieval conferences, emphasizing the *innovative technological aspects* of search and retrieval systems.

It gathers researchers and practitioners from both *academia and industry* to discuss the latest innovative and visionary ideas in the field.


still 16 oz.
still 16 oz.
still 16 oz.
still 16 oz.
still 16 oz.
Vermont's Finest
BEN & JERRY'S
ICE CREAM
half baked
Chocolate & Vanilla Ice Creams
mixed with Gobs of Chocolate Chip
Cookie Dough & Fudge Brownies

DESIRES

A systems-oriented biennial conference, complementary in its mission to the mainstream Information Access and Retrieval conferences, emphasizing the *innovative technological aspects* of search and retrieval systems.

It gathers researchers and practitioners from both *academia and industry* to discuss the latest innovative and visionary ideas in the field.

Source: Flickr (michael_hendrickx/5327363552)

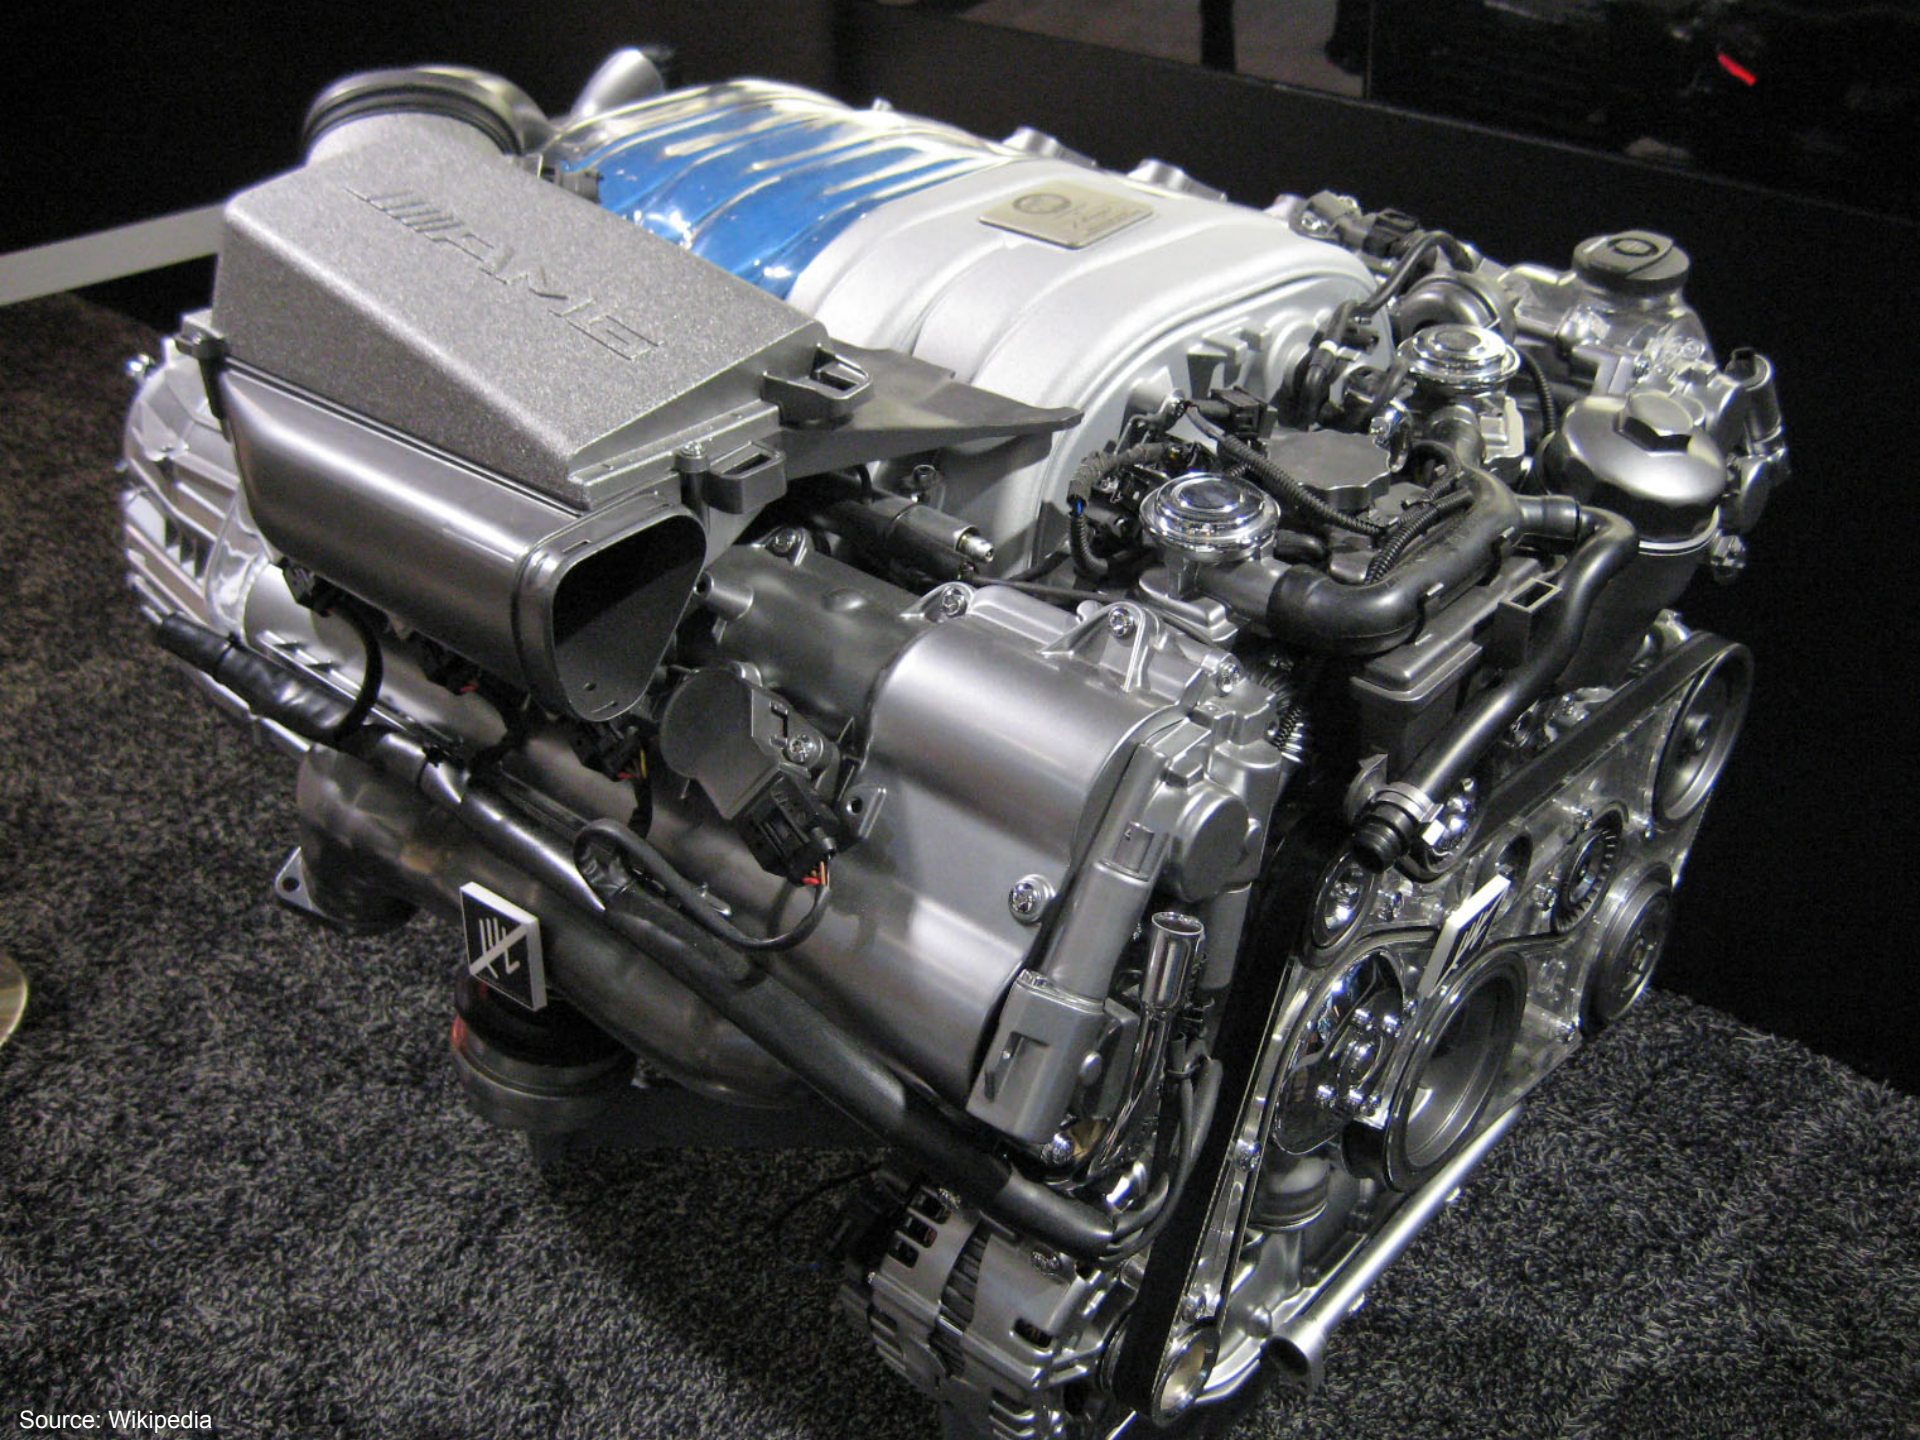

Source: Wikipedia

Source: https://qz.com/1013634/spacex-launch-watch-live-as-elon-musks-rocket-company-launches-bulgarias-first-communications-satellite/

What's the connection?

Source: Flickr (agrusoft/26223084353/)

Source: Google


Source: Google

Source: Flickr (massimoleonardi/8064194094/)

Source: Flickr (nnova/5110654026/)

Source: Flickr (claudiu92/5392298355/)

Source: Flickr (72712524@N04/8225583084/)

Source: Flickr (massimoleonardi/8064194094/)

Source: Flickr (massimoleonardi/8064194094/)

Cloud computing allows us to explore different abstractions and organizations of computing

In the beginning…

Source: Wikipedia

Infrastructure as a Service (IaaS)

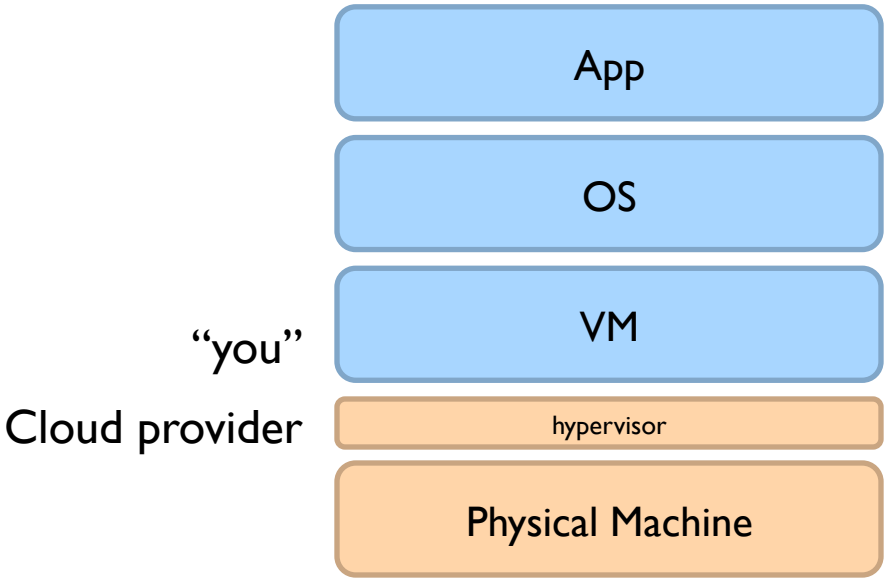

Multi-Tenancy

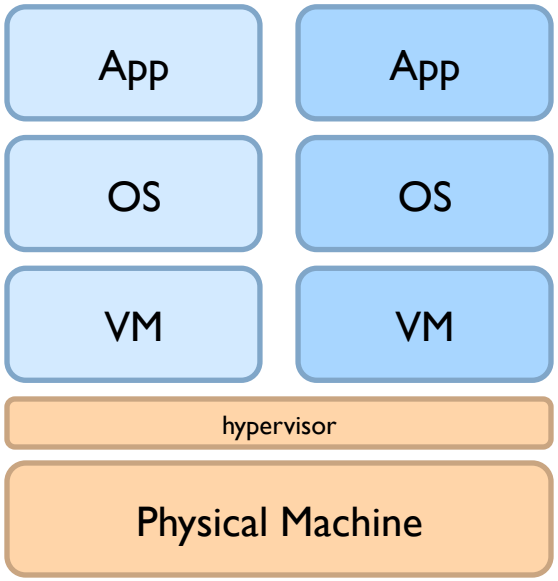

Containers >> VMs

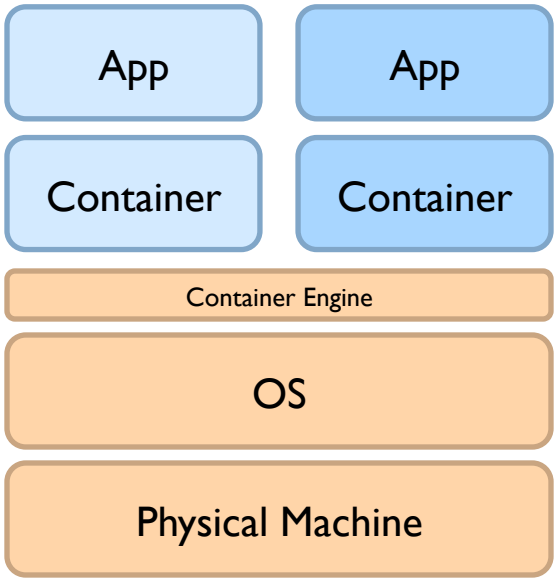

Typical Stack

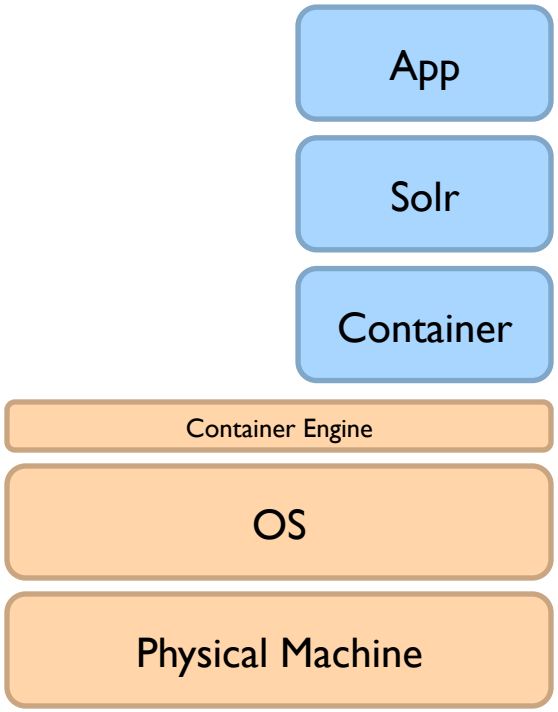

Multi-Container Orchestration

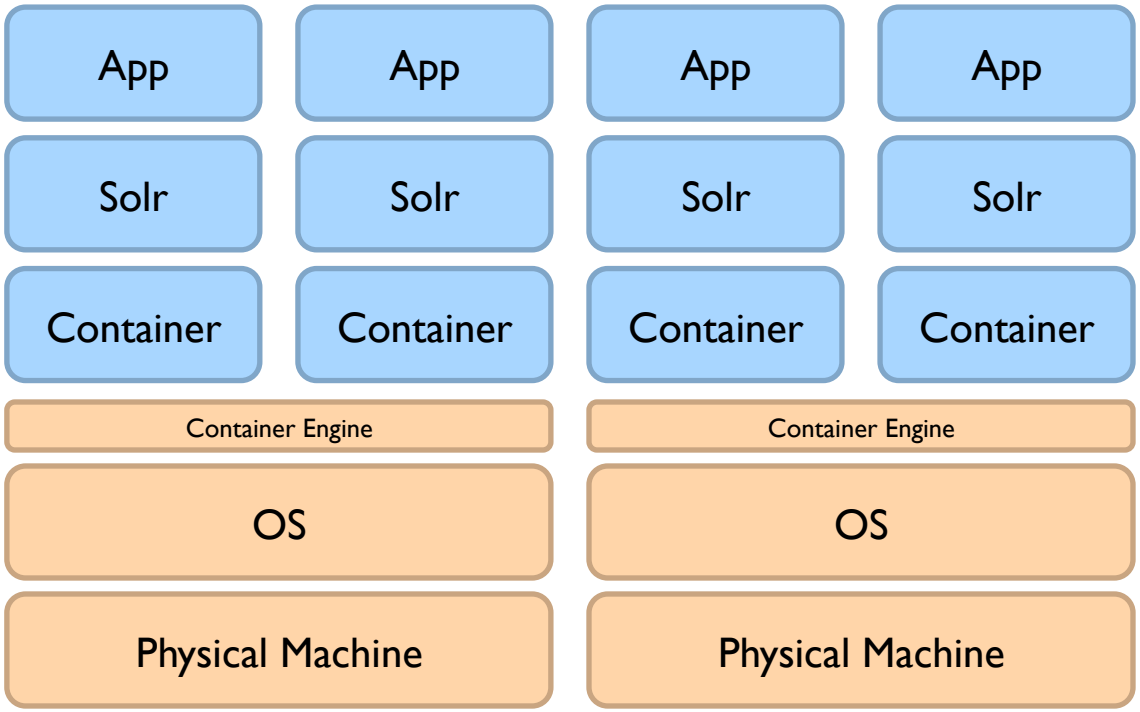

Source: Flickr (purplelight/14754940092/)

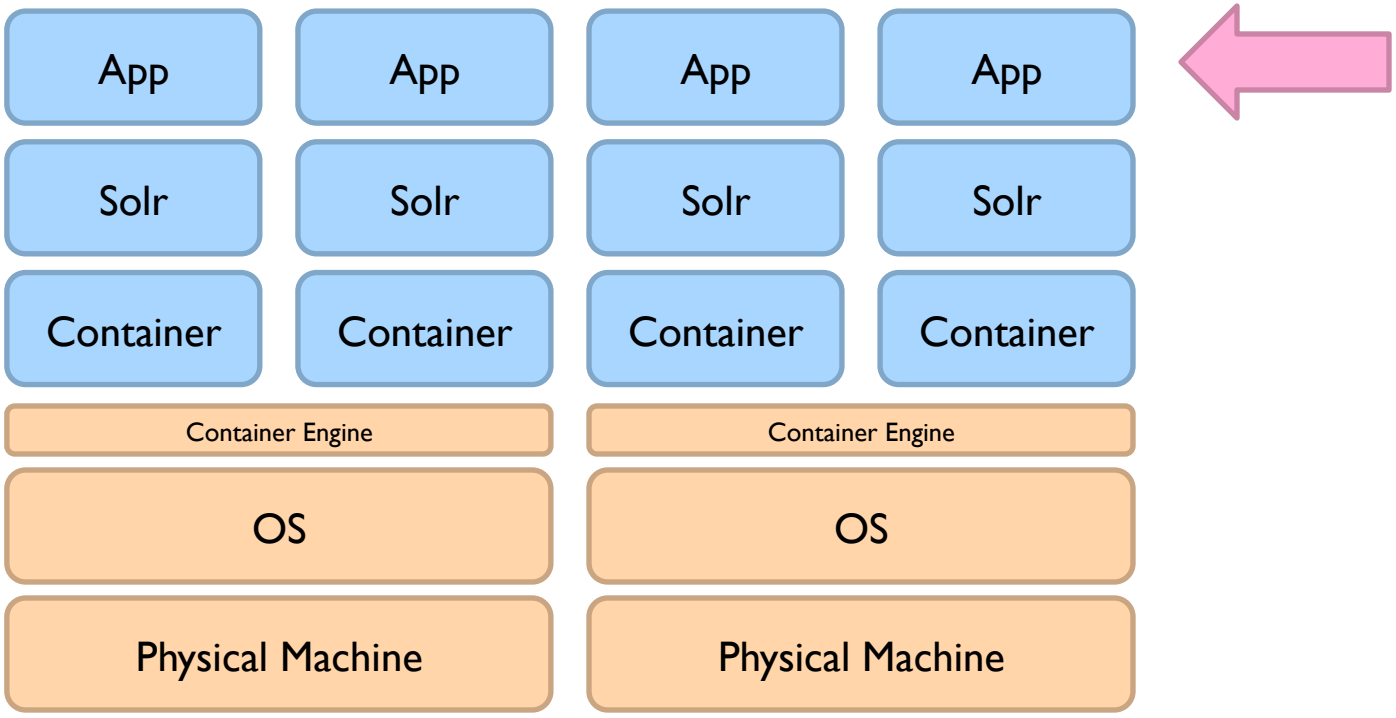
App
App
App
App
Solr
Solr
Solr
Solr
Container
Container
Container
Container
Container Engine
Container Engine
OS
OS
Physical Machine
Physical Machine

Platform as a Service

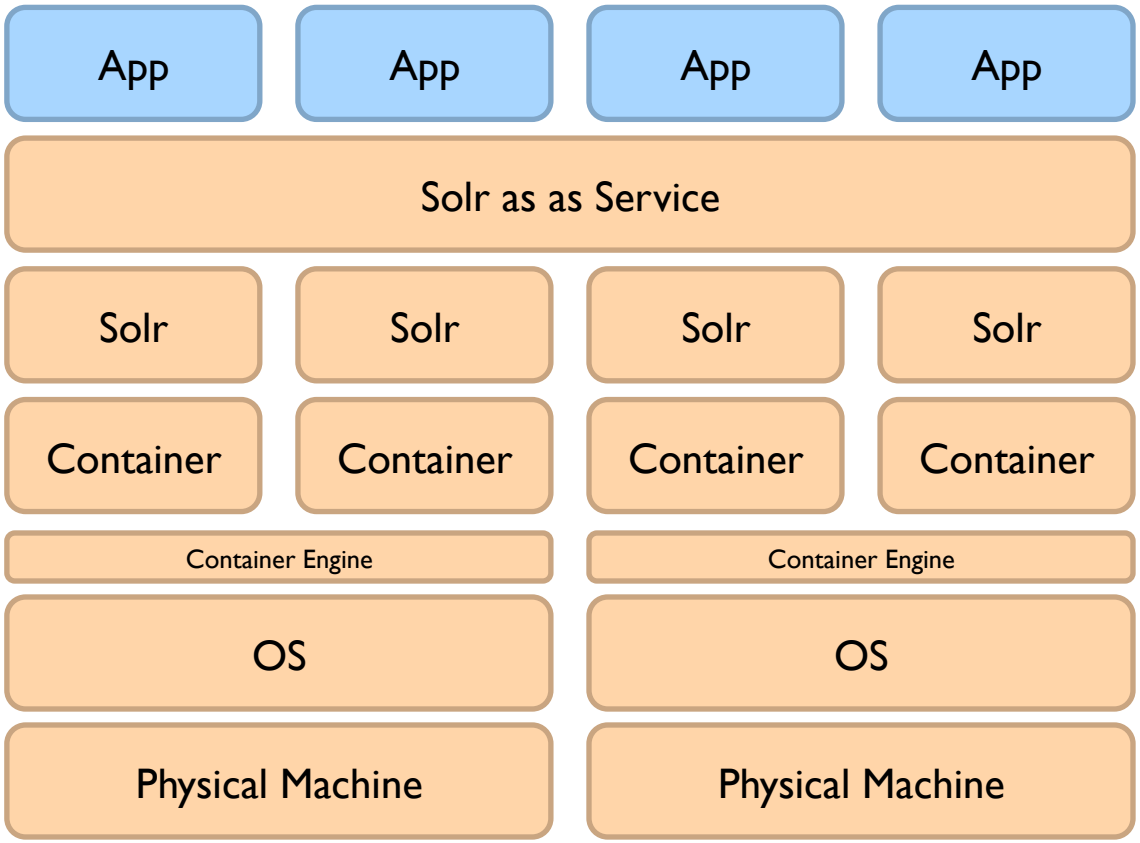

Platform as a Service

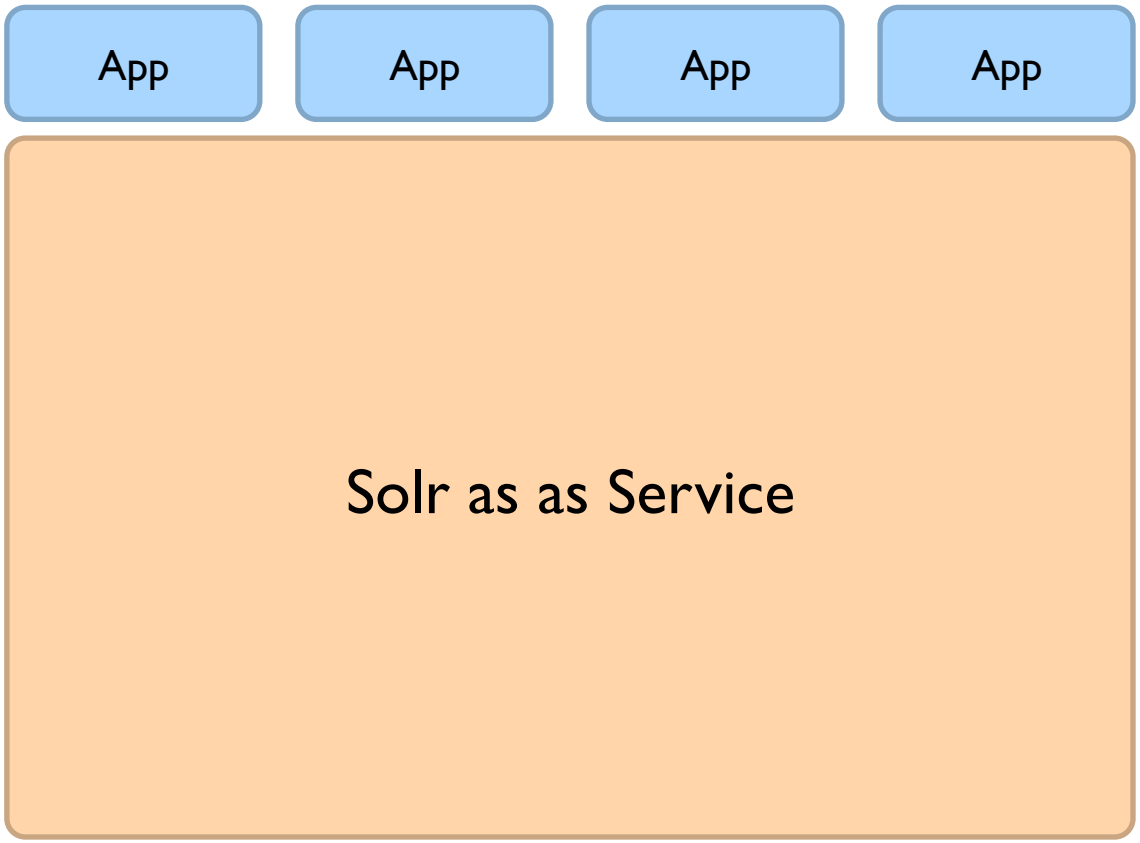

Service-level agreements (SLA) in terms
of latency, capacity, scalability, etc.

What about the apps?

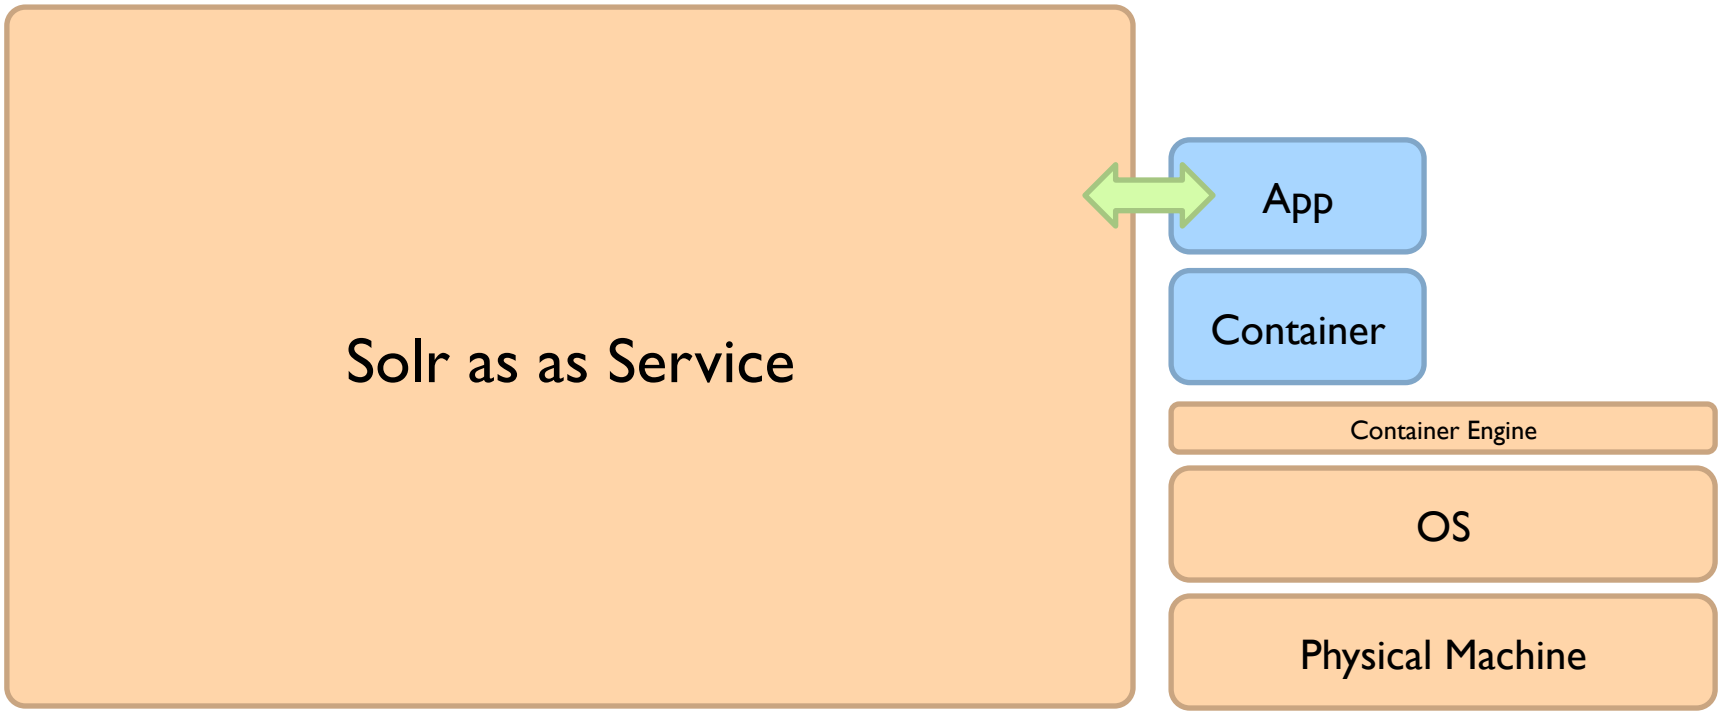

Scaling out the apps…

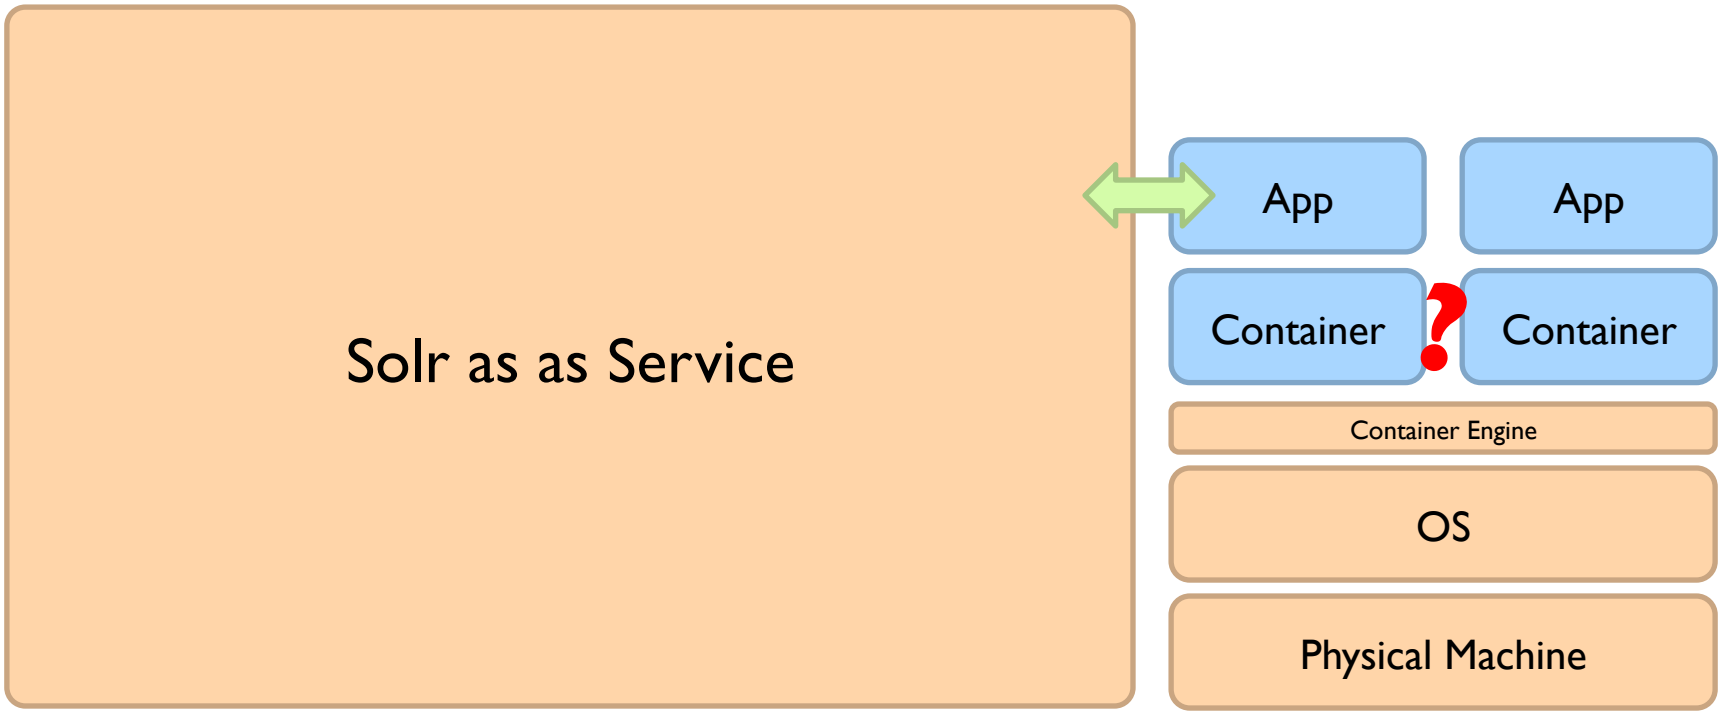

Source: Flickr (purplelight/14754940092/)

Source: Flickr (michael_hendrickx/5327363552)

Operational Semantics of Computing…

$$\frac{\langle E, s\rangle \Rightarrow V}{\langle L := E\,,\ s\rangle \longrightarrow (s \uplus (L \mapsto V))}$$

if the expression E in state s reduces to value V,
then the program $L := E$ will update the state s with the assignment $L = V$

State	State as a Service (persistent storage, databases, message queues, …)
Transitions	Function as a Service (blocks of code with a well-defined entry and exit points)

Source: Wikipedia

Computing without Servers

Developer: Write a bunch of functions

typically – read state, perform some computation, update state

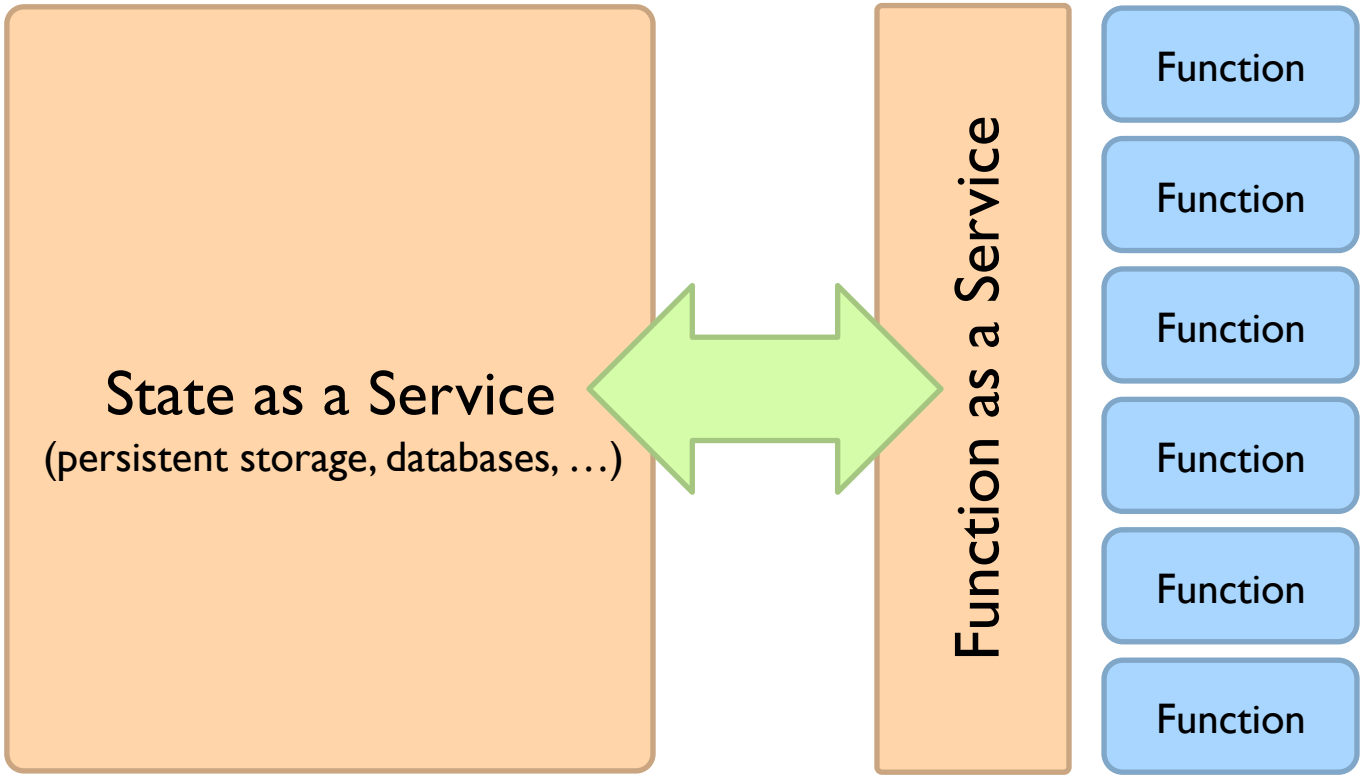

Cloud provider handles everything else!

allocation of resources for execution, scaling up and down, load balancing, cleaning up, etc.

Cost model: pay per function invocation

Serverless Examples

Example: *Image Thumbnail Creation*

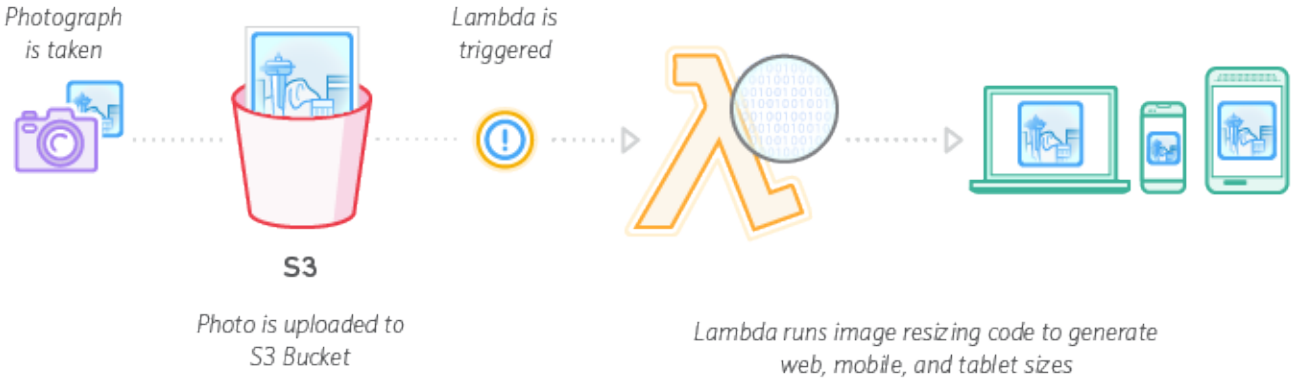

Example: *Retail Data Warehouse ETL*

Source: Amazon Web Services

Serverless Examples

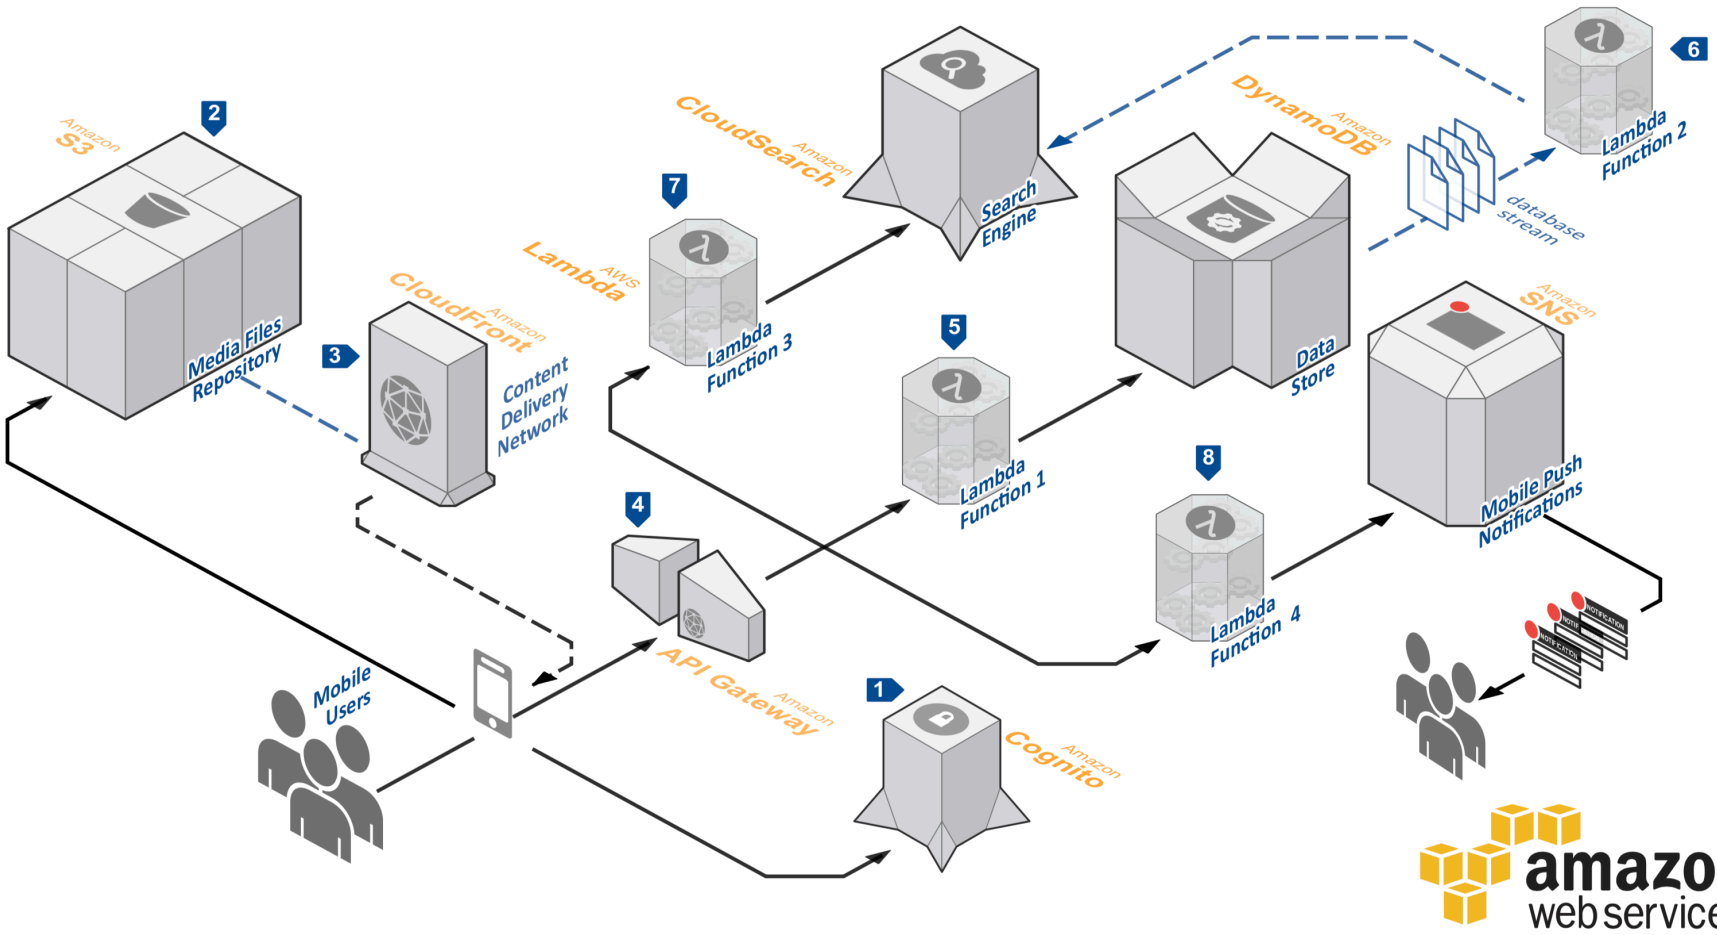

Source: Amazon Web Services

Source: Flickr (michael_hendrickx/5327363552)

It's not *actually* computing without servers…

It's just someone else's problem!

What would a serverless search engine look like?

Serverless Search

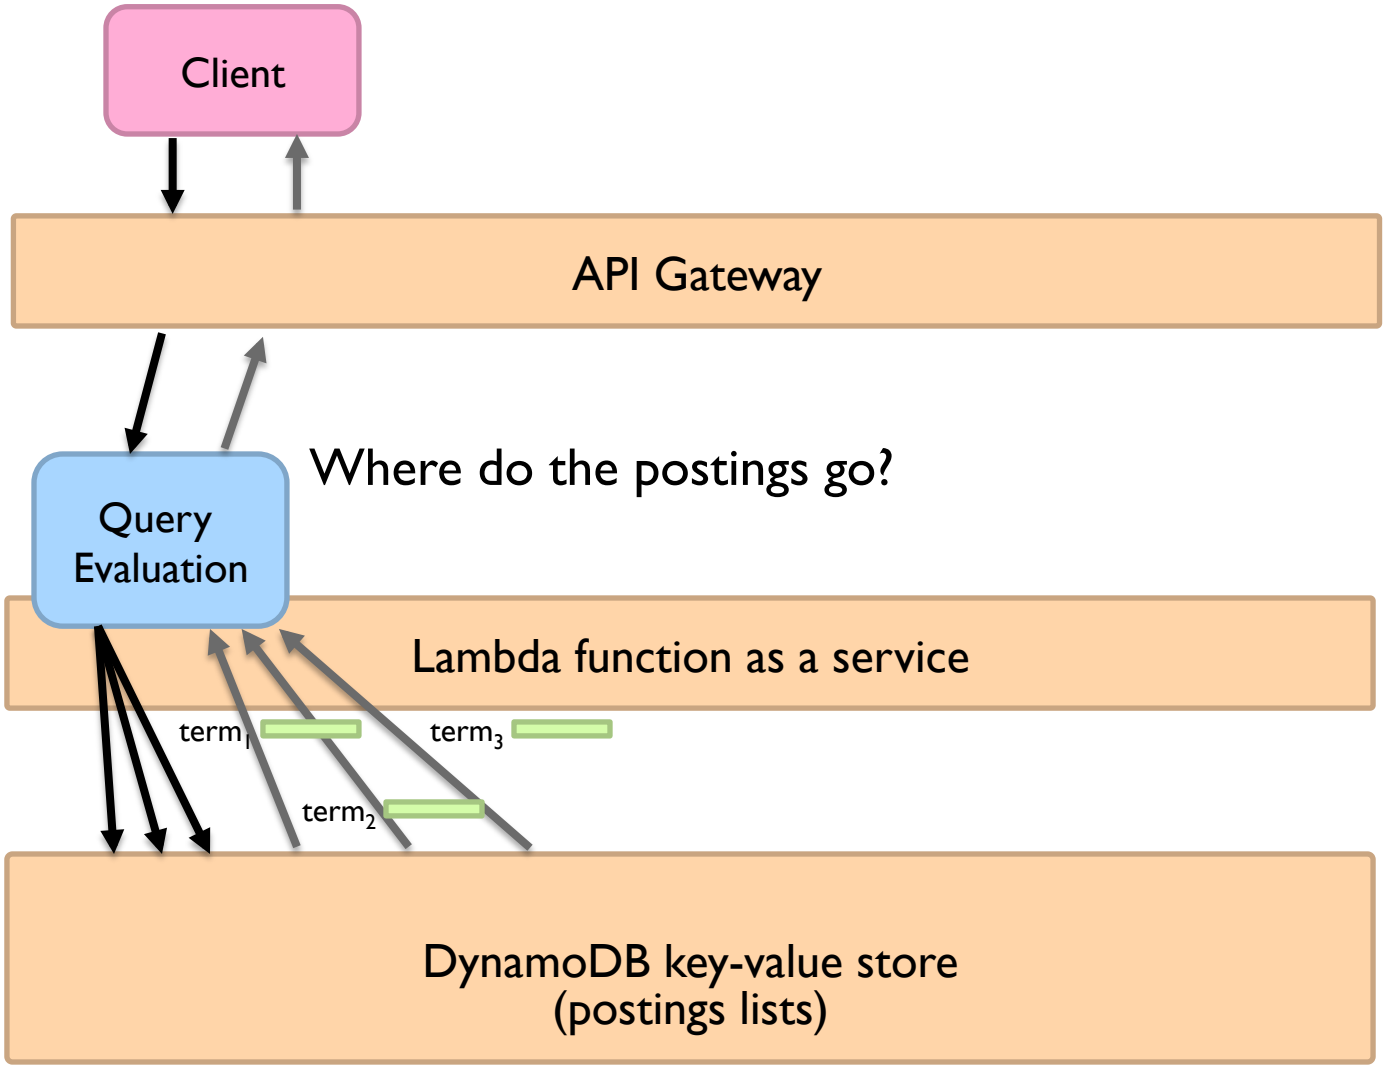

Matt Crane and Jimmy Lin. An Exploration of Serverless Architectures for Information Retrieval. *ICTIR 2017.*

Serverless Search

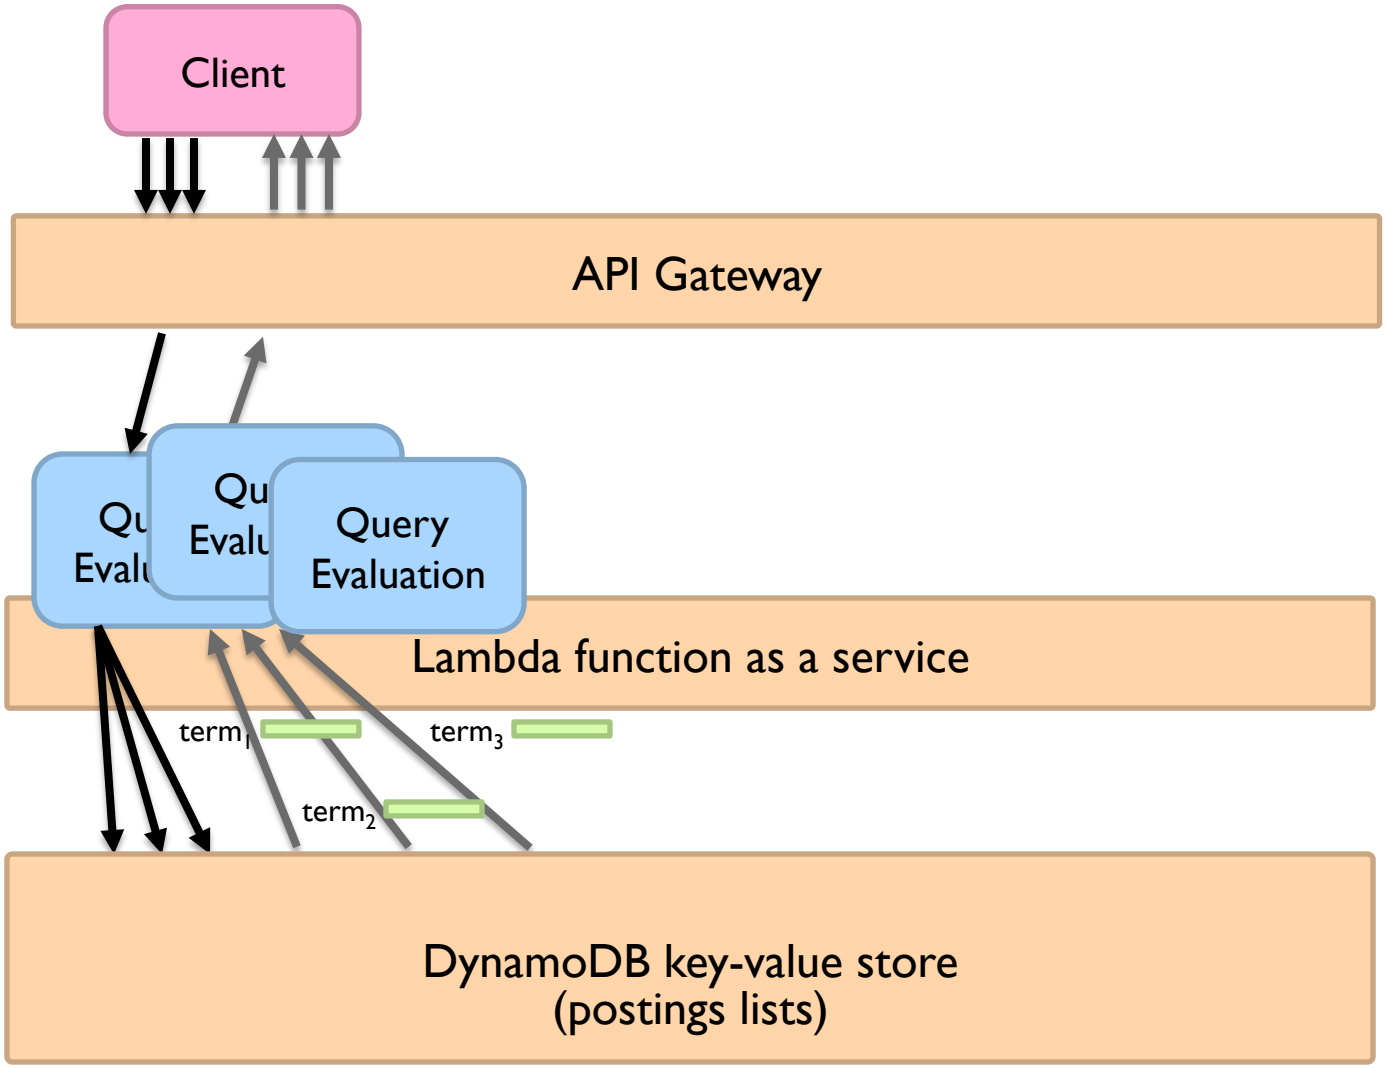

Matt Crane and Jimmy Lin. An Exploration of Serverless Architectures for Information Retrieval. *ICTIR 2017.*

Serverless Search

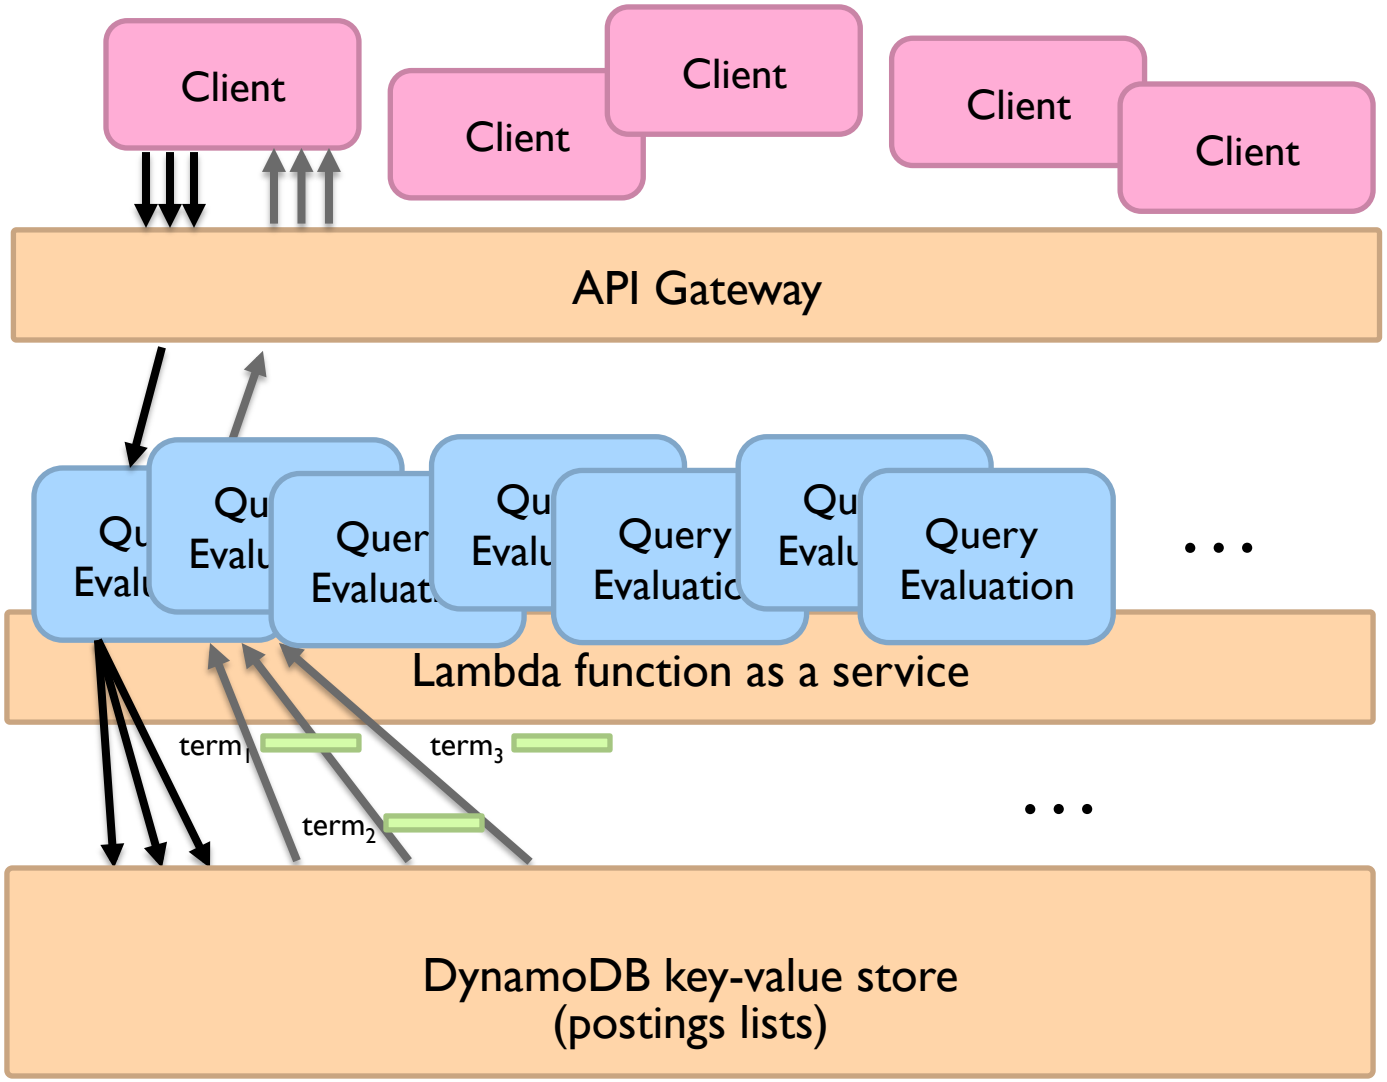

Matt Crane and Jimmy Lin. An Exploration of Serverless Architectures for Information Retrieval. *ICTIR 2017.*

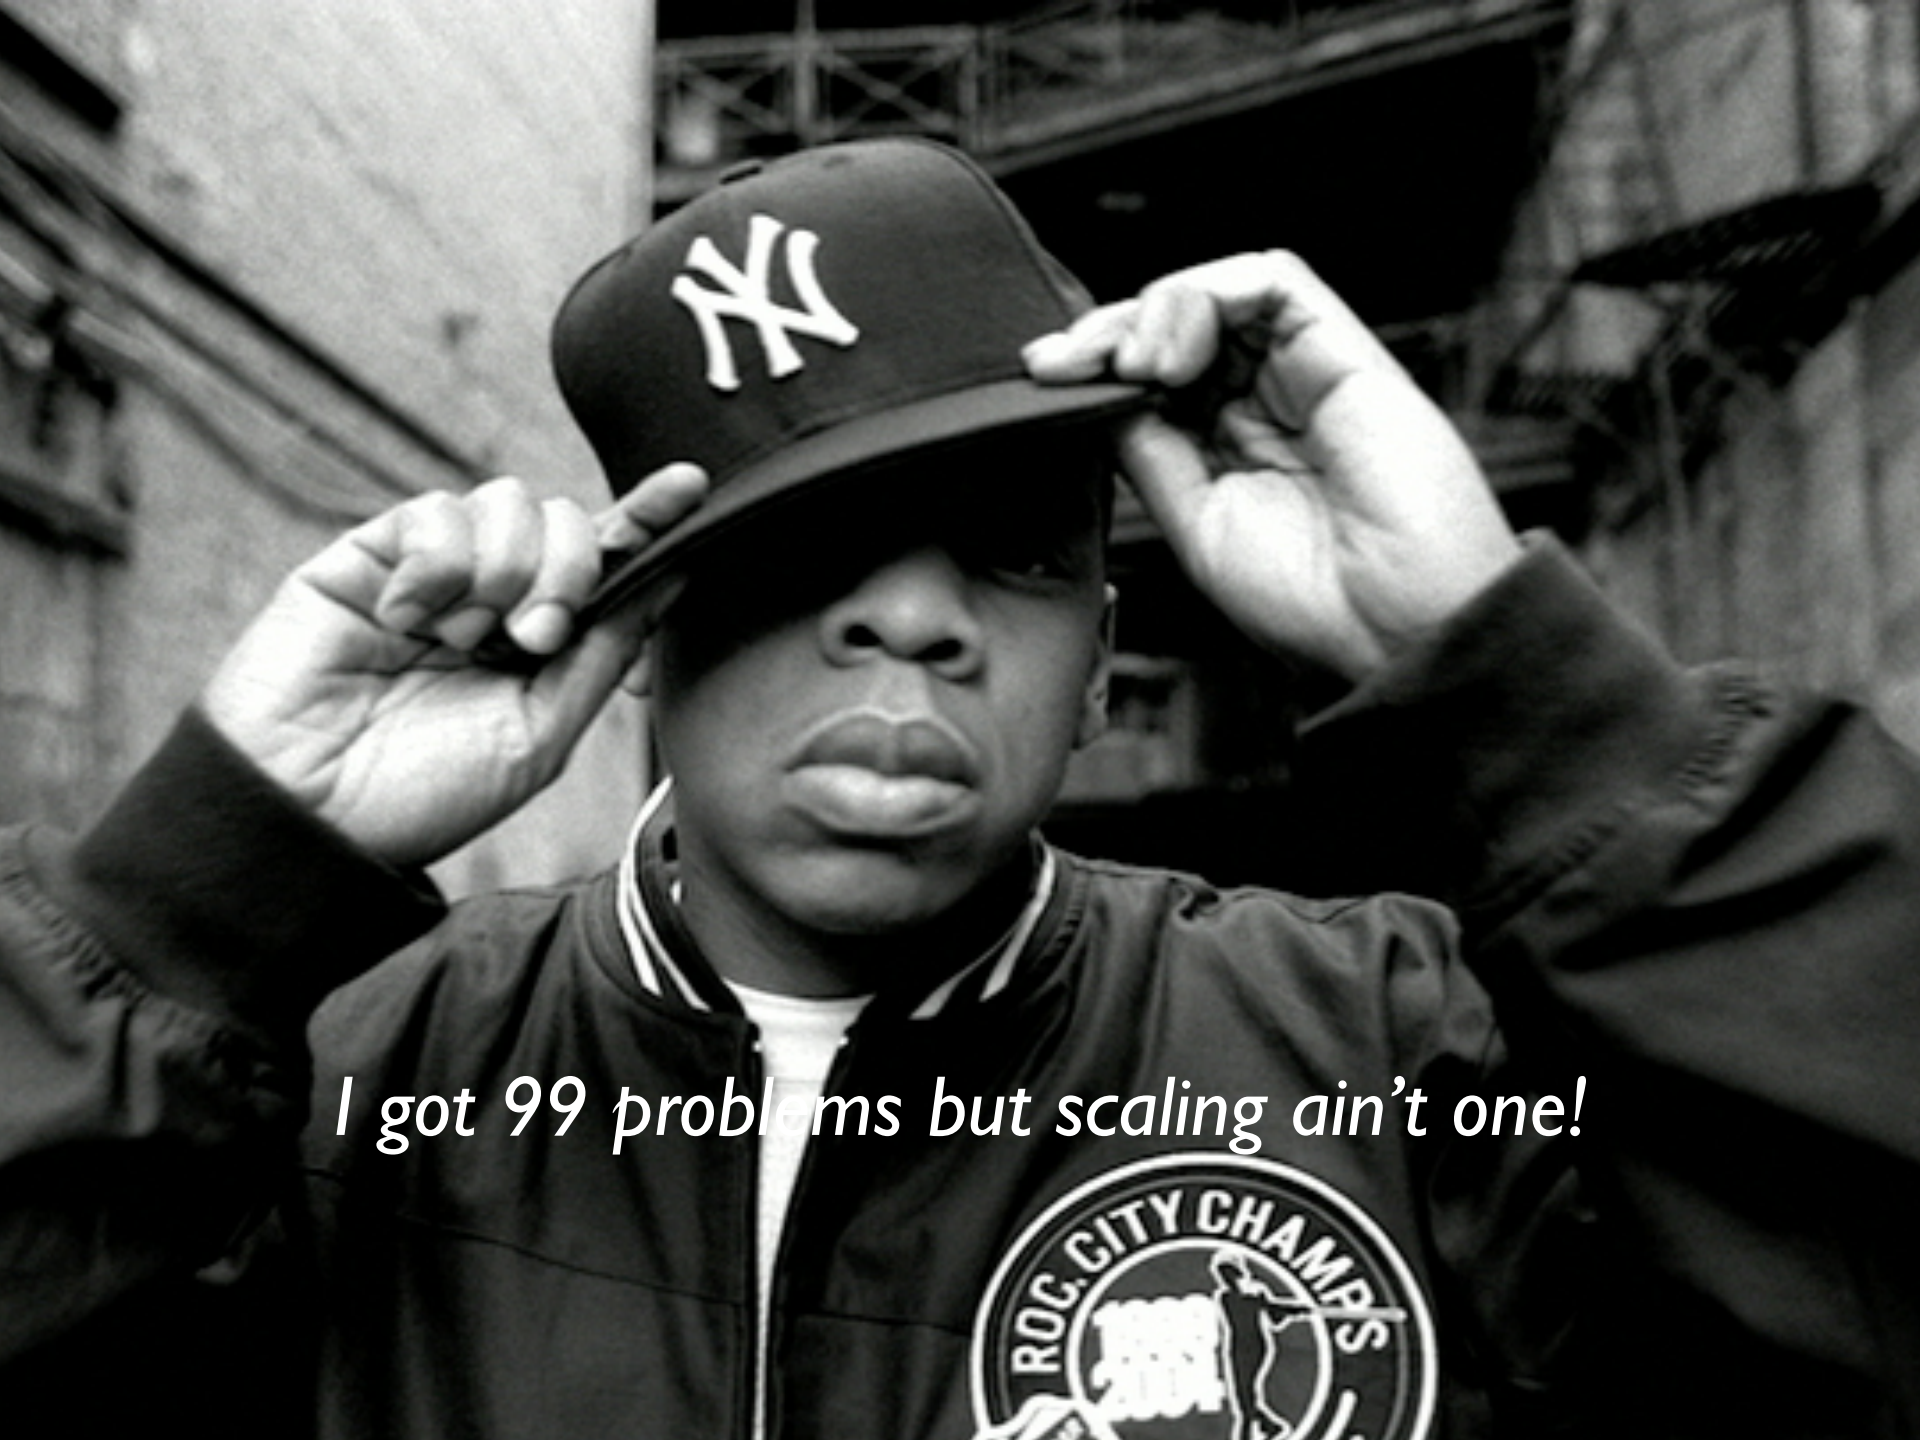
I got 99 probl ms but scaling ain't one!
ROC CITY CHAMPS

How well does serverless search work?

tl; dr – *currently, not very well…*

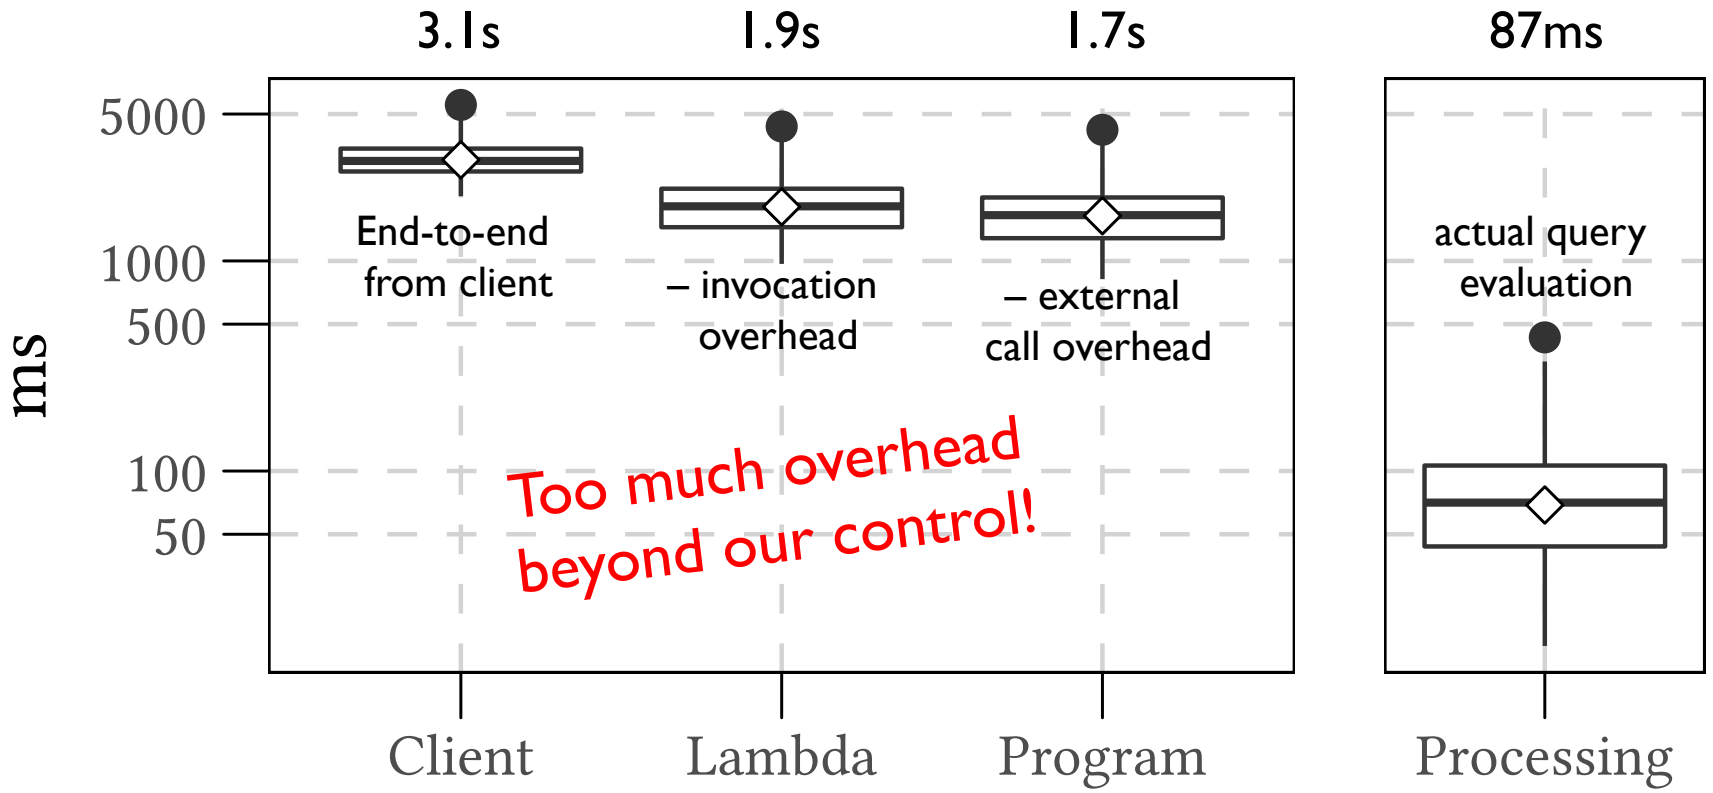

Query Latency

Gov2 collection (25m docs), topic 701-850

Matt Crane and Jimmy Lin. An Exploration of Serverless Architectures for Information Retrieval. *ICTIR 2017.*

Serverless Search Costs

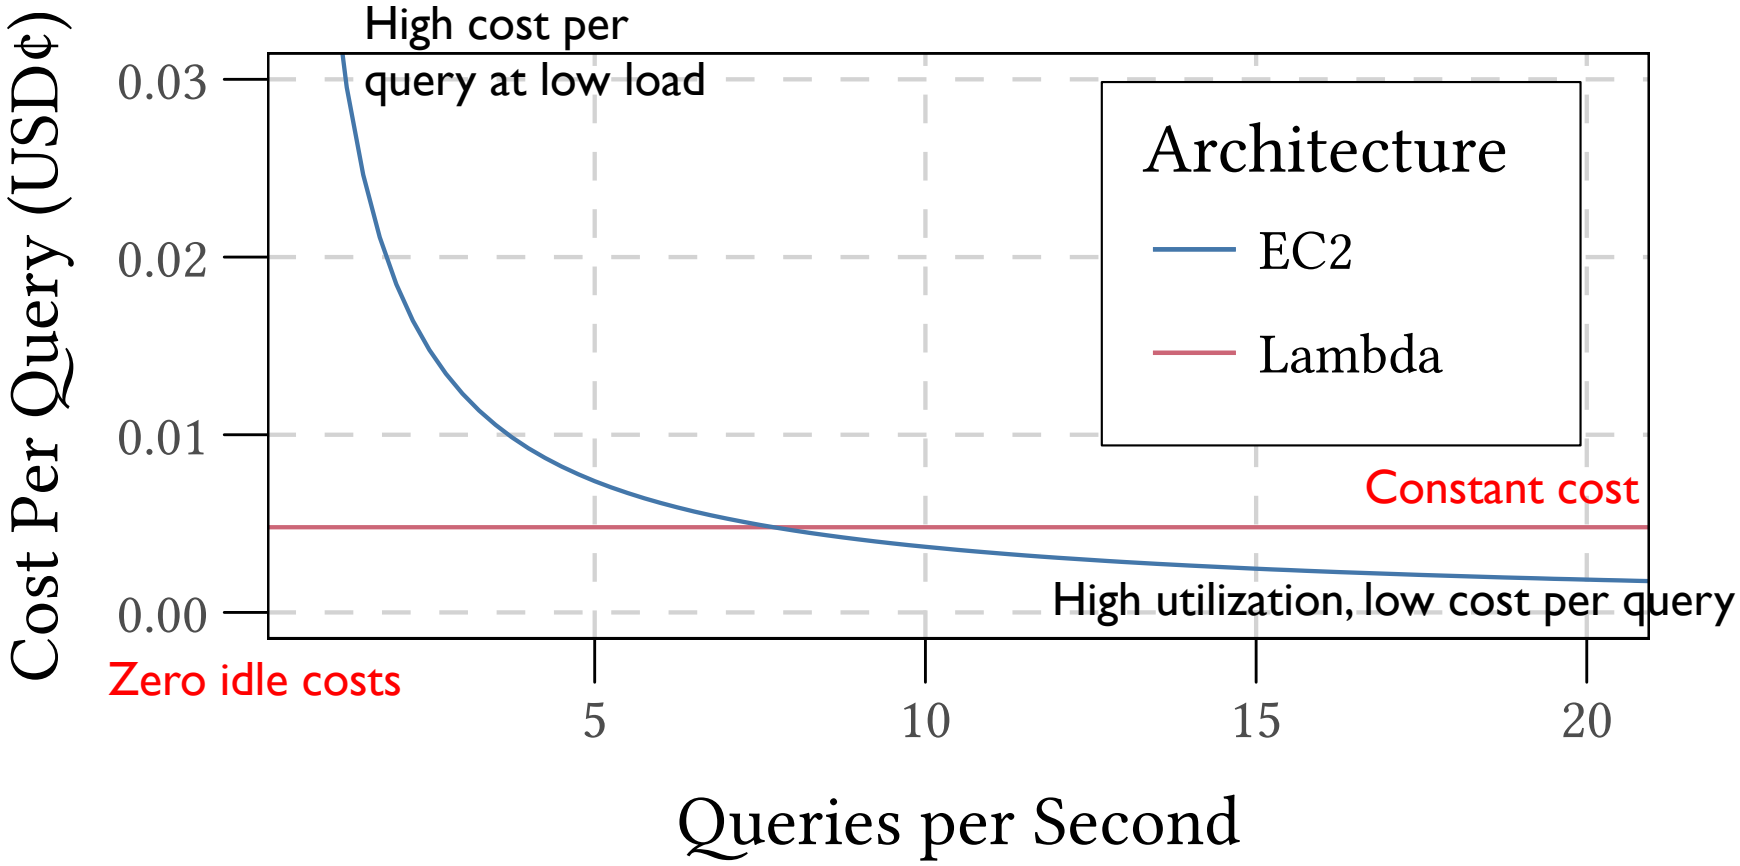

EC2 r3.4xlarge instance – extrapolating from JASS
(ECIR 2016 reproducibility study)

Matt Crane and Jimmy Lin. An Exploration of Serverless Architectures for Information Retrieval. *ICTIR 2017.*

Okay, what about neural networks?

Serverless Neural Networks

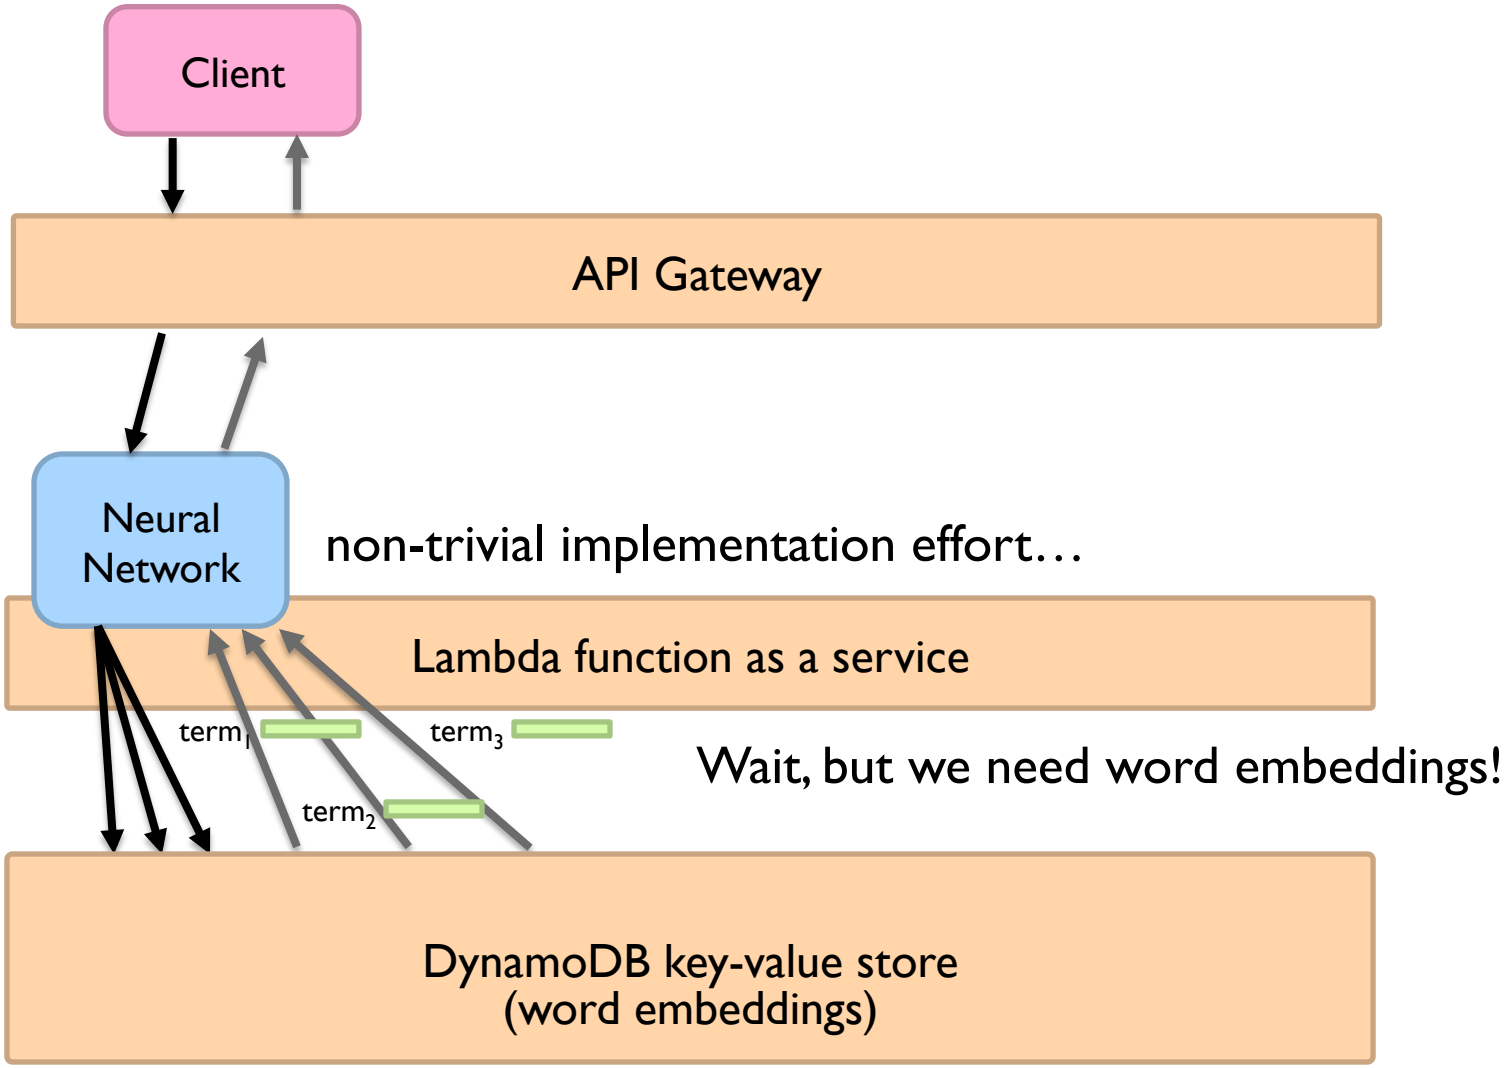

Zhucheng Tu, Mengping Li, and Jimmy Lin. Pay-Per-Request Deployment of Neural Network Models Using Serverless Architectures. *NAACL 2018 Demo.*

How well does it work?

tl; dr – *pretty compelling!*

Neural Network for answer selection in question answering
(Severyn and Moschitti, SIGIR 2015)

C	tput	Latency (ms)			Cost
	(QPS)	mean	p50	p99	(/10^6 Q)
5	12.1	410	381	657	$1.04
10	21.1	468	443	780	$1.04
15	30.8	467	439	827	$1.04
20	38.5	496	486	785	$1.04
25	44.4	530	519	814	$1.25

3.8 million queries per day
$4.80 USD per day

Zhucheng Tu, Mengping Li, and Jimmy Lin. Pay-Per-Request Deployment of Neural Network Models Using Serverless Architectures. *NAACL 2018 Demo.*

More than embarrassingly parallel?

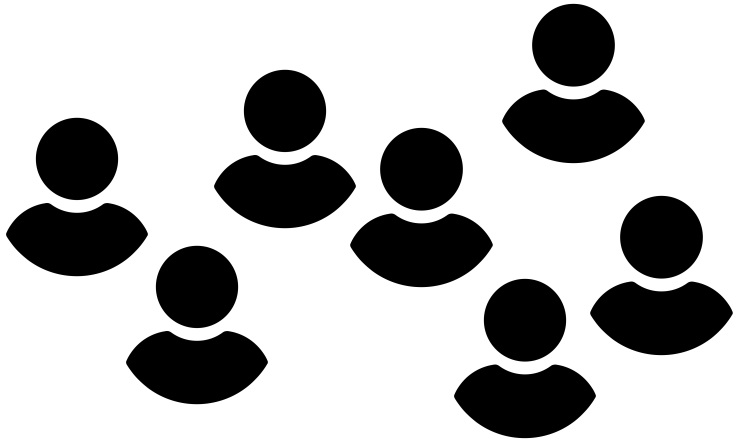

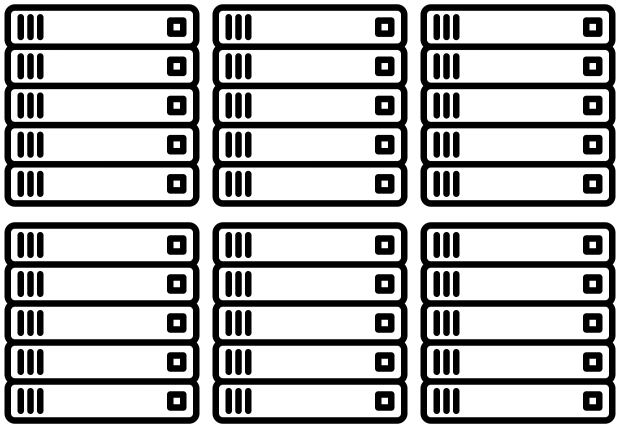

Managing clusters is a pain!

Source: Flickr (michael_hendrickx/5327363552)

Flint is Serverless (Py)Spark

Implementation of Spark cluster backend using AWS Lambda

Seamless experience:
point configs at Flint, fire up PySpark shell, go!

Youngbin Kim and Jimmy Lin. Serverless Data Analytics with Flint. *Workshop on Serverless Computing (WoSC 2018)*.

Flint is Serverless PySpark!

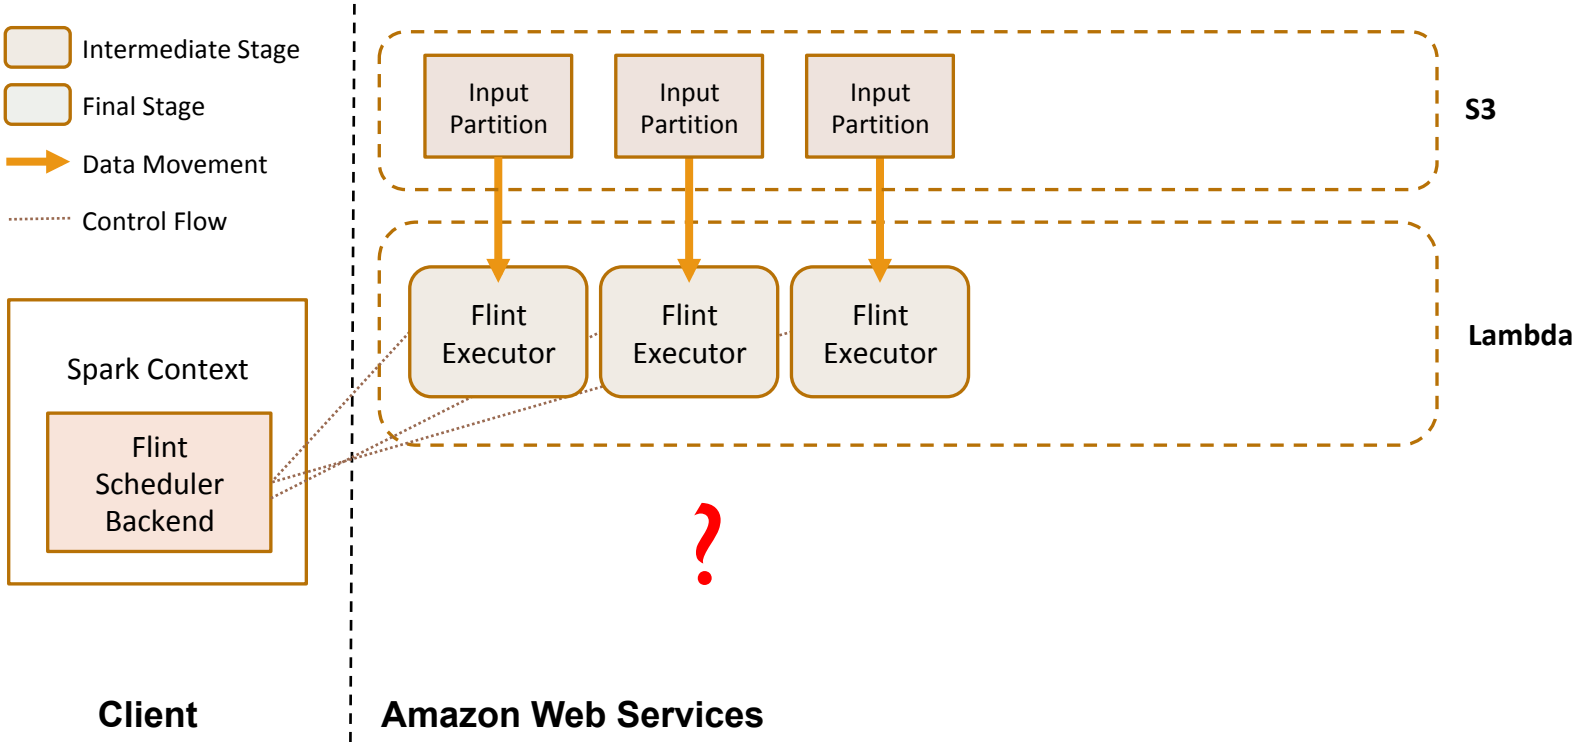

Youngbin Kim and Jimmy Lin. Serverless Data Analytics with Flint. *Workshop on Serverless Computing (WoSC 2018).*

What about data shuffling?

Obvious solution: write to S3
Downside: need to implement our own shuffling mechanism

Flint's (clever) idea: use Amazon's queuing service (SQS)
Map each group to a queue – AWS handles data movement

Data shuffling as a service!

Flint is Serverless PySpark!

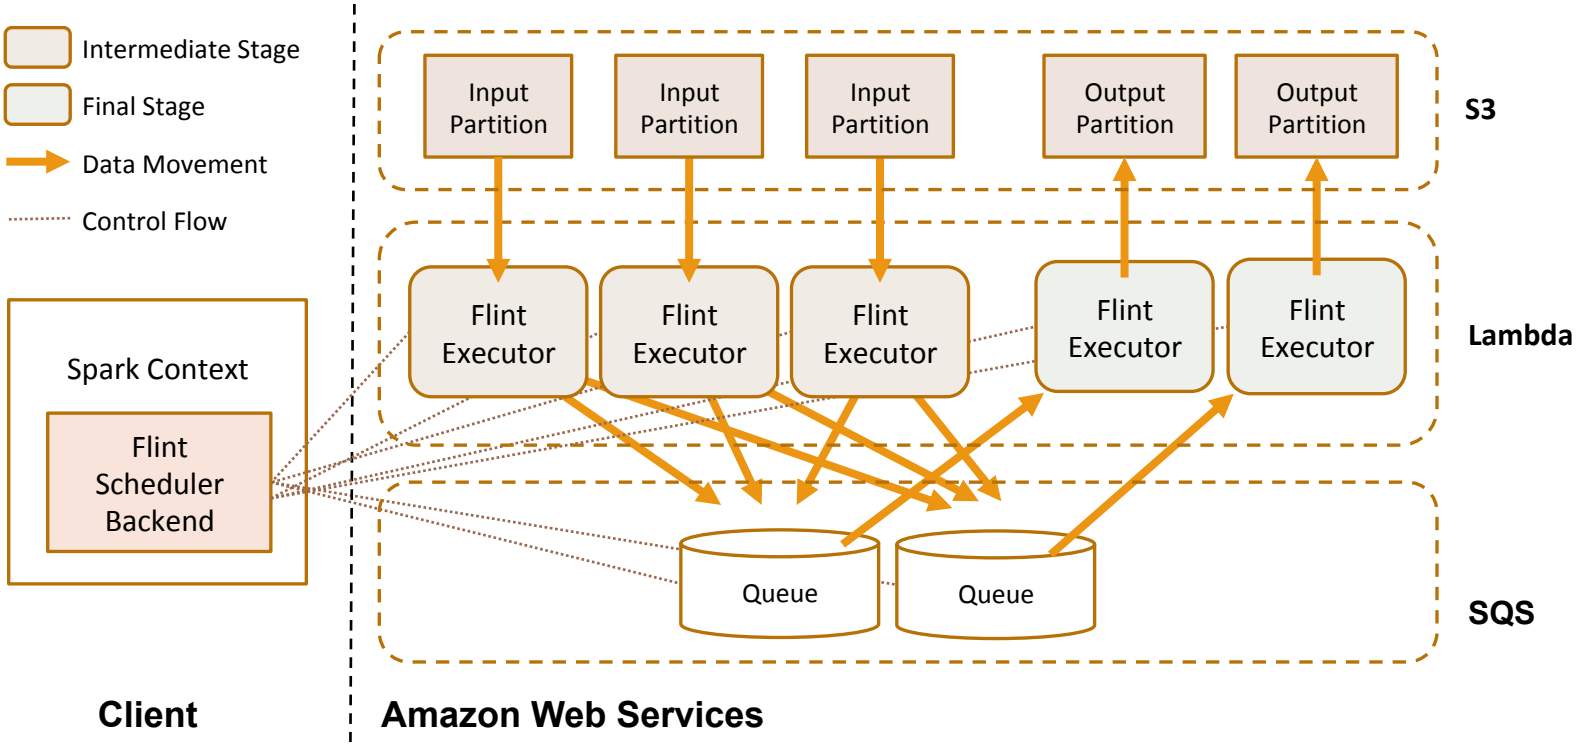

Youngbin Kim and Jimmy Lin. Serverless Data Analytics with Flint. *Workshop on Serverless Computing (WoSC 2018).*

How well does it work?

215 GB dataset, 80 concurrent Lambda executors

Simple filter/aggregation queries

Performance and cost are on par with Spark!

A couple of minutes, ~$0.50 USD

Youngbin Kim and Jimmy Lin. Serverless Data Analytics with Flint. *Workshop on Serverless Computing (WoSC 2018).*

Source: Flickr (michael_hendrickx/5327363552)

The Balance of Computing

Where does computing happen?

Smart

Dumb

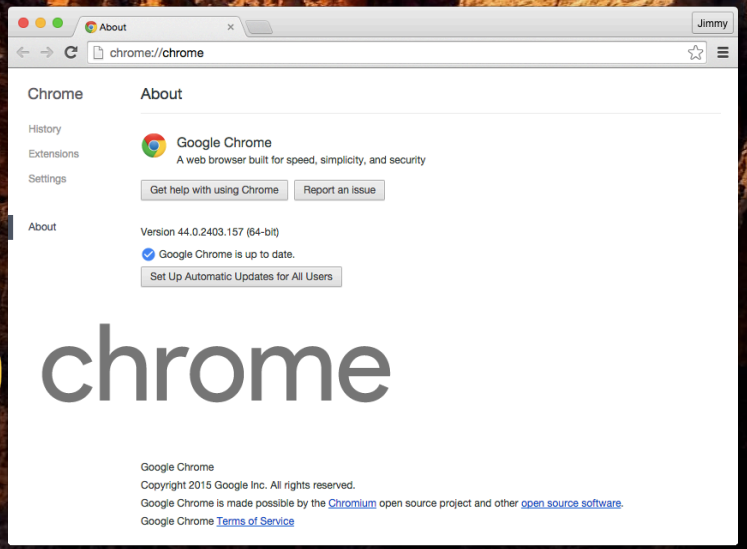

Source: Flickr (massimoleonardi/8064194094/)

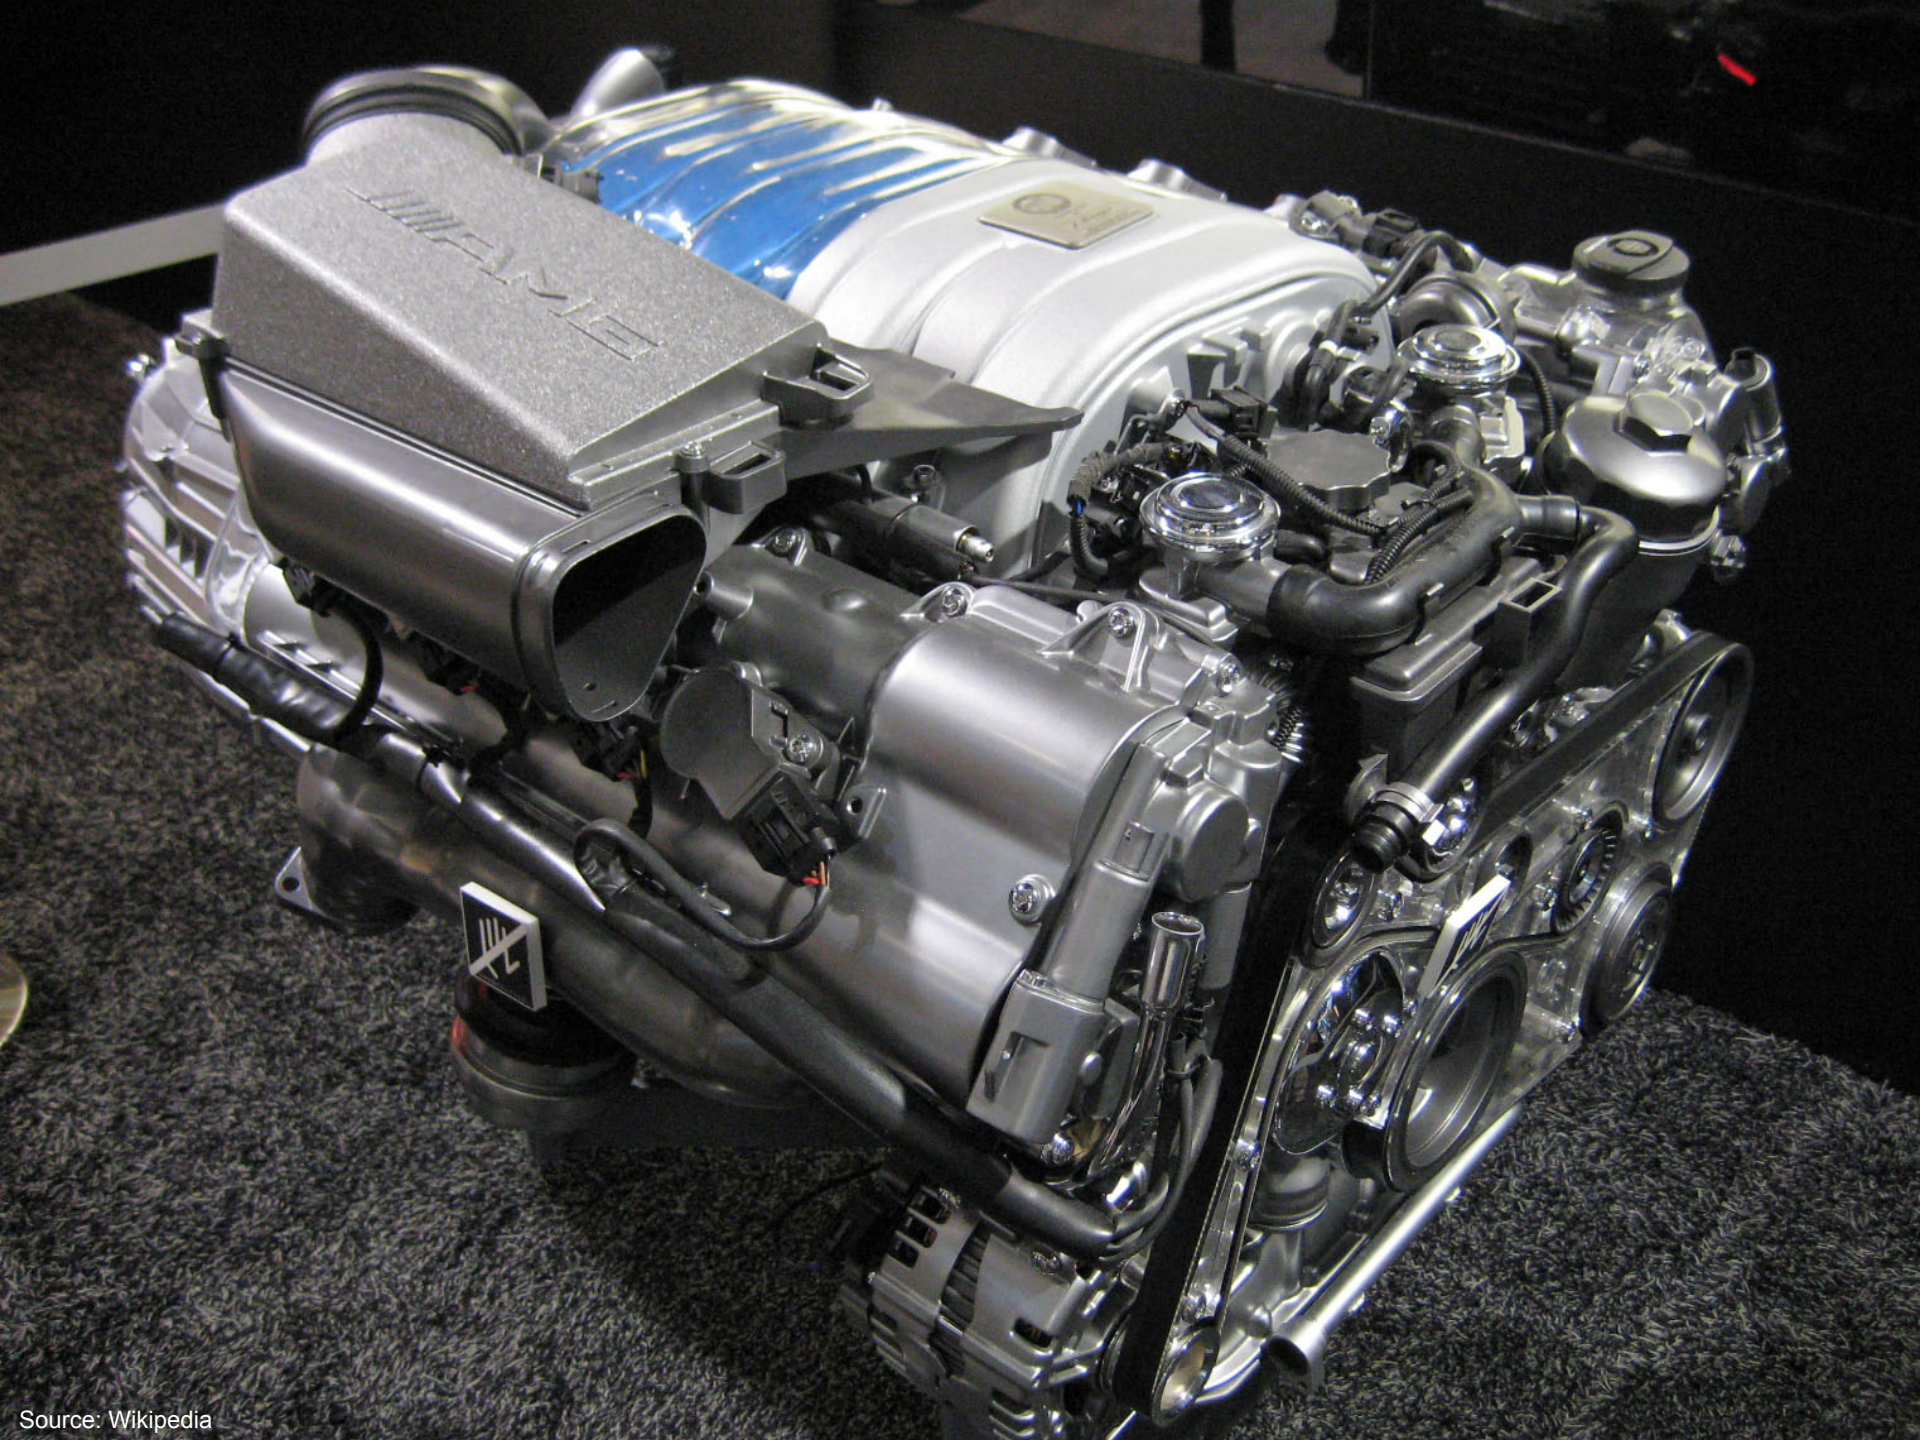

Source: Wikipedia

No viable alternative to JavaScript in the browser…

(better make it fast!)

Google's high performance, open source, JavaScript engine.

Might as well run it server side also…

(Yay microservices!)

Incredibly powerful combination!


Free For All
Waiting for Players
60fps
0 : 2 3
50
10
5
10
50
5
50
5
0ups
30
0
50
100
100

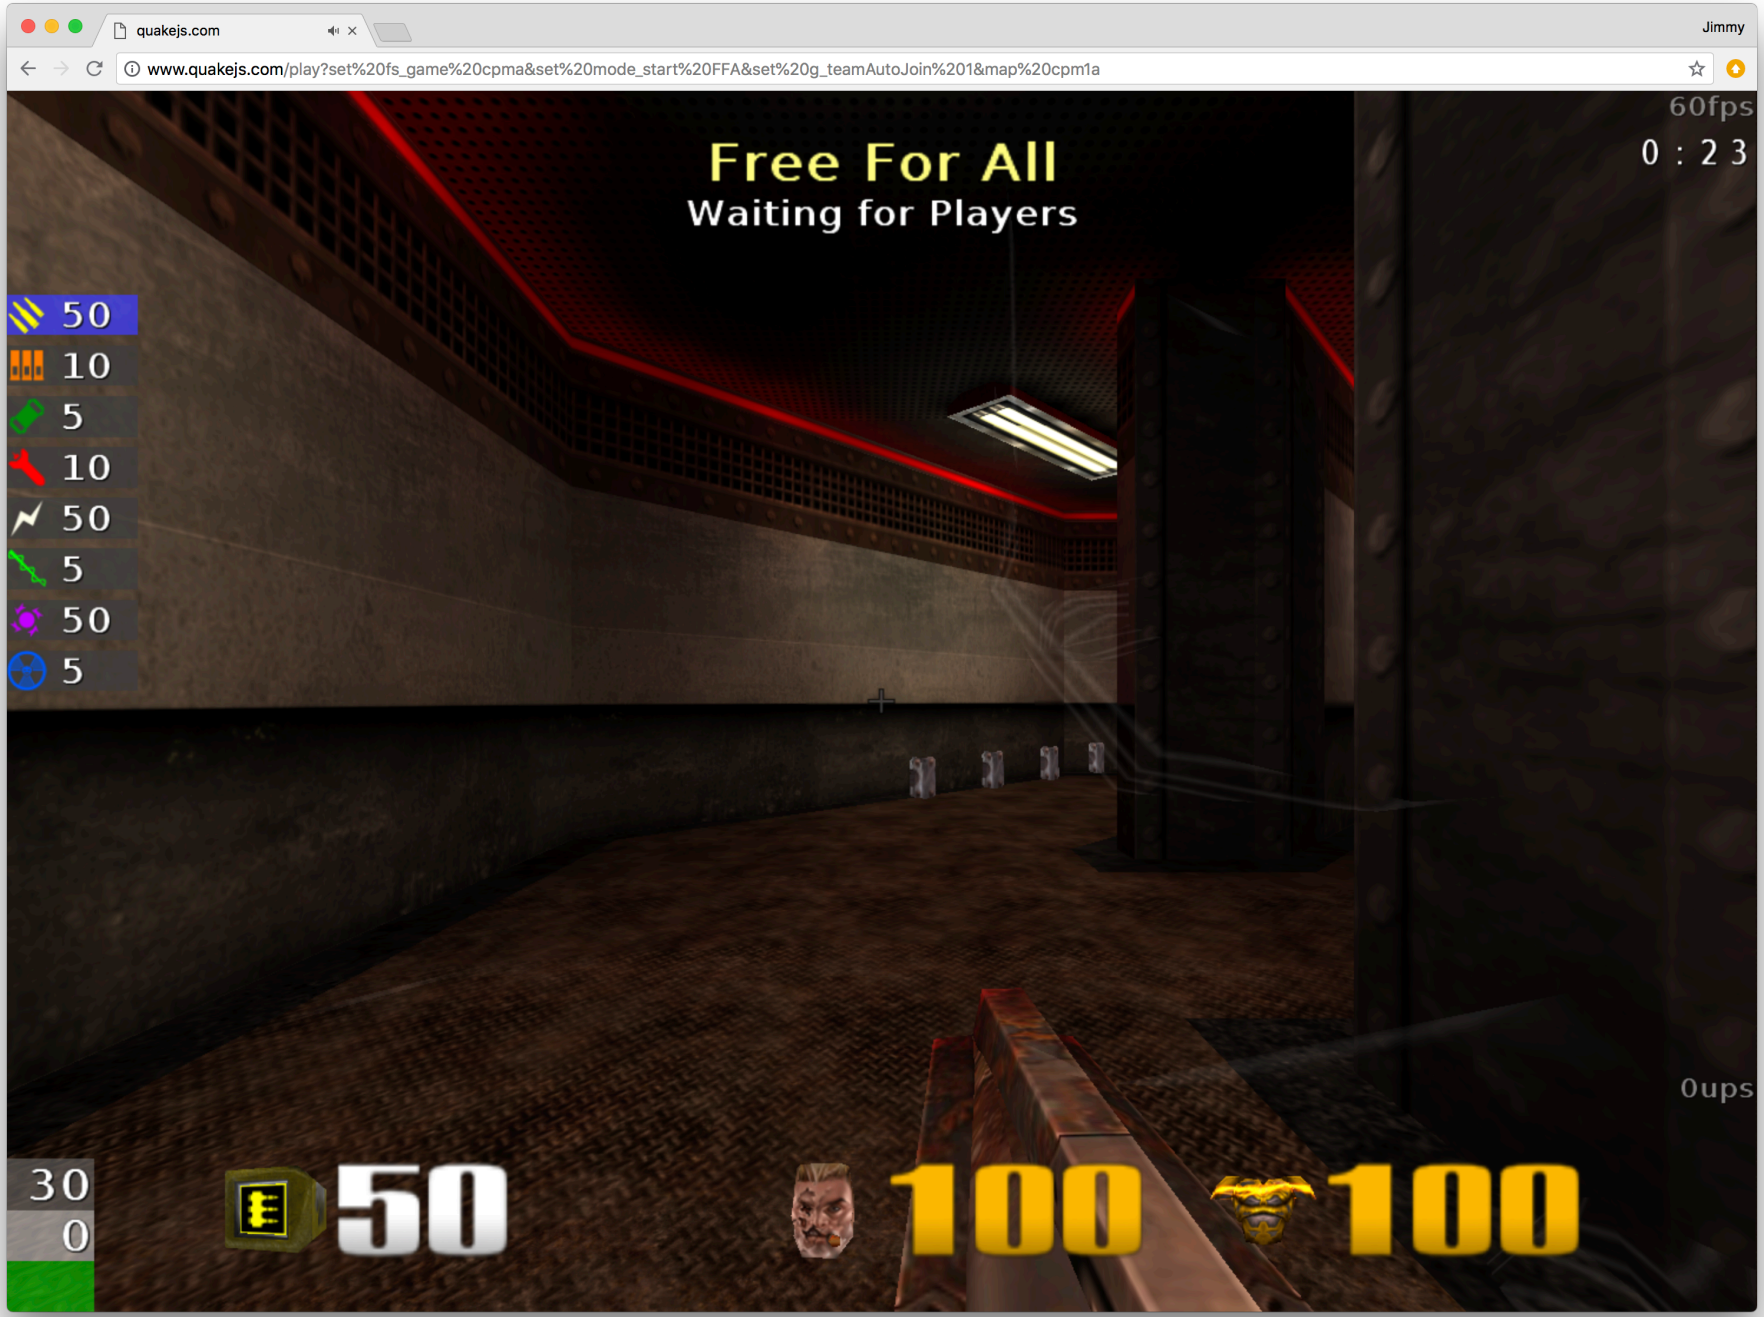
quakejs.com
Jimmy
www.quakejs.com/play?set%20fs_game%20cpma&set%20mode_start%20FFA&set%20g_teamAutoJoin%201&map%20cpm1a
Free For All
Waiting for Players
60fps
0 : 2 3
50
10
5
10
50
5
50
5
0ups
30
0
50
100
100

Source: Wikipedia

What would an in-browser JavaScript search engine look like?

A Self-Contained In-Browser Search Engine

IndexDB for storing postings

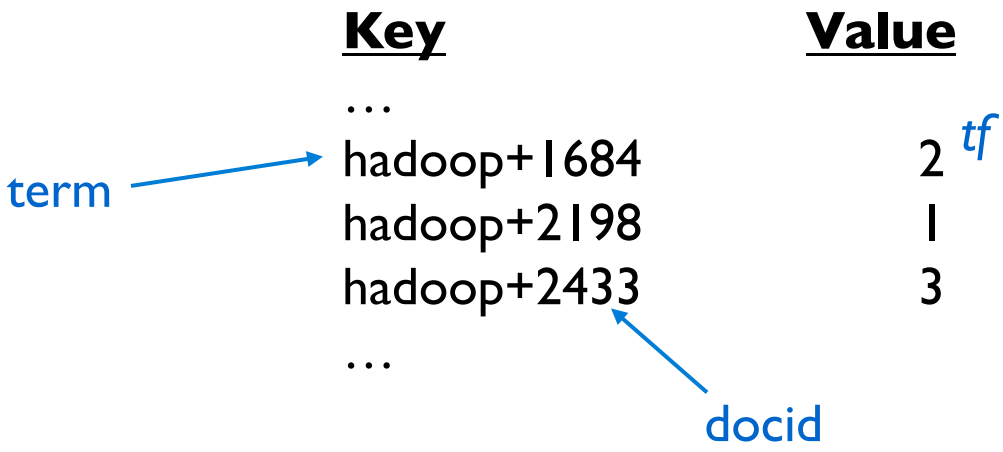

Term-at-a-Time query evaluation in JavaScript

(natural fit with async dispatches in JavaScript)

Jimmy Lin. Building a Self-Contained Search Engine in the Browser. *ICTIR 2015.*

How well does it work?

tl; dr – *performance kinda sucks*

Collection: 1.12m tweets; Queries: TREC 2011/2012 MB Track
Experiments on a 2012 Macbook Pro (Chrome browser)

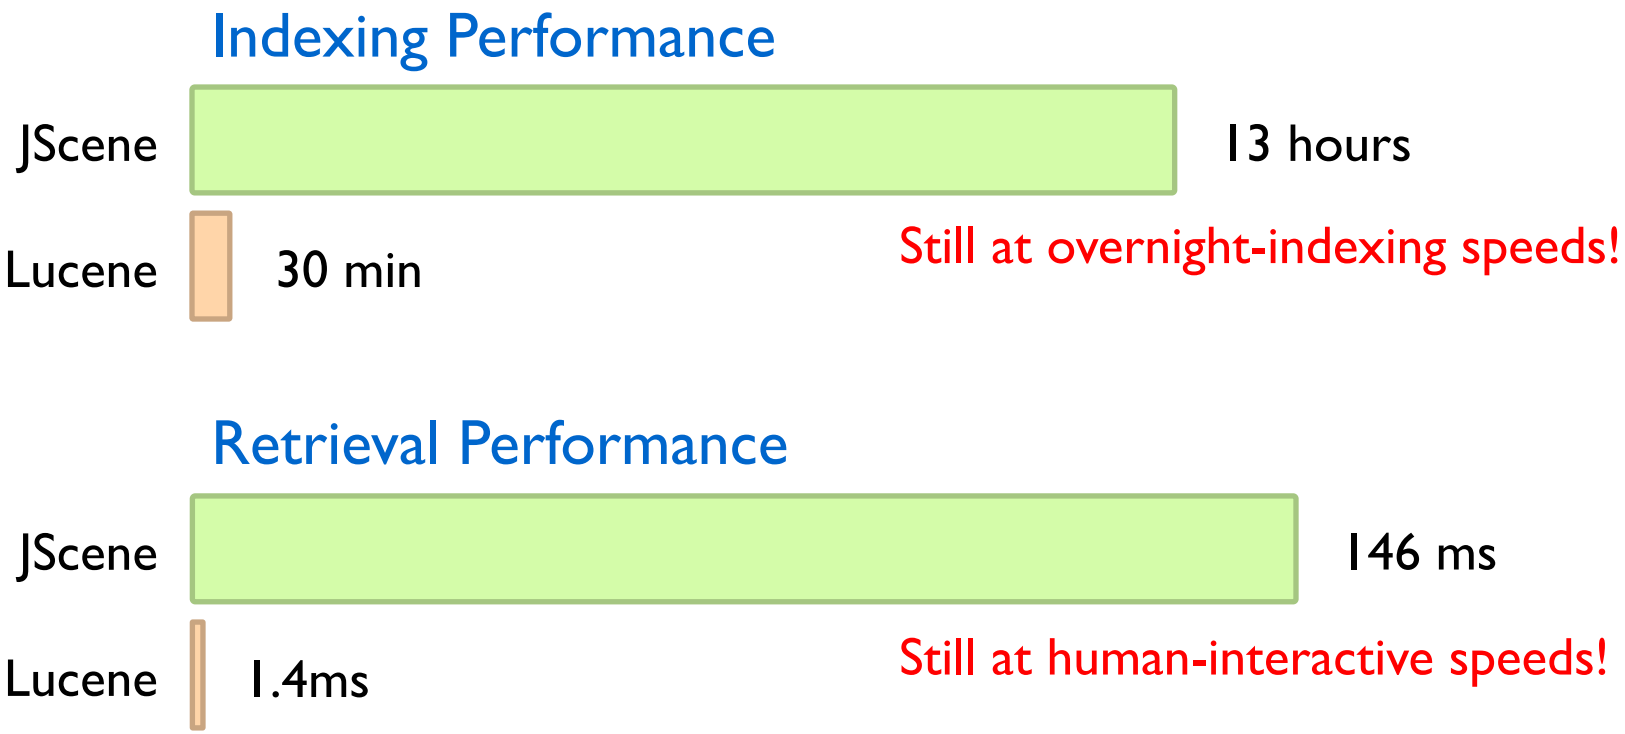

Jimmy Lin. Building a Self-Contained Search Engine in the Browser. *ICTIR 2015.*

What would an in-browser JavaScript search engine look like?

What about a database?

Kareem El Gebaly and Jimmy Lin. In-Browser Interactive SQL Analytics with Afterburner. *SIGMOD 2017 Demo.*

Afterburner

Interactive SQL analytics in your browser!

A SQL query engine in JavaScript, embeddable in any webpage…

Trick #1: columnar data layout w/ typed arrays

Trick #2: query compilation into asm.js

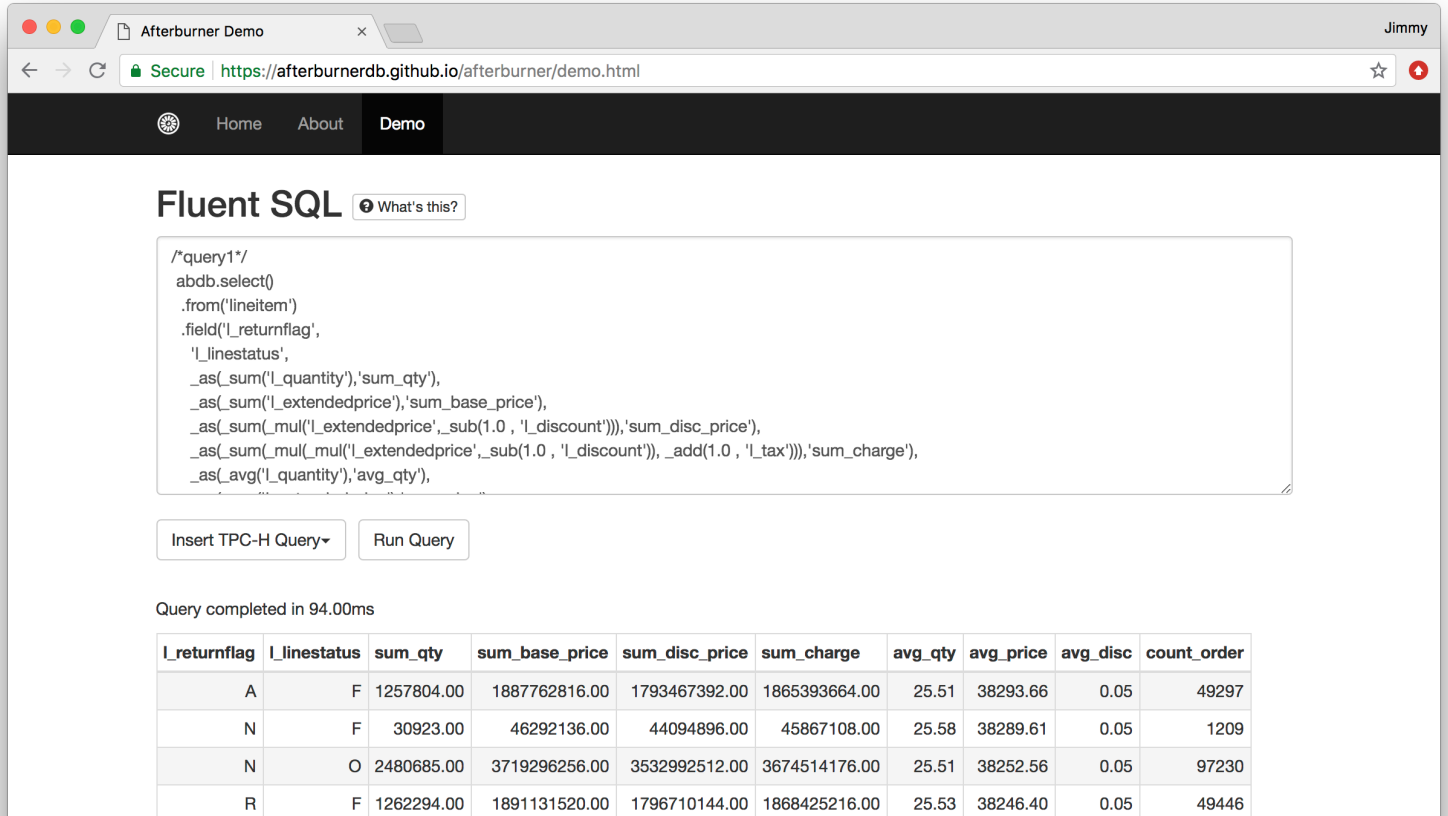

How well does it work?

tl; dr – *performance competitive with sota query processing!*

TPC-H, scale factor 1 GB = 6m rows in lineitem table

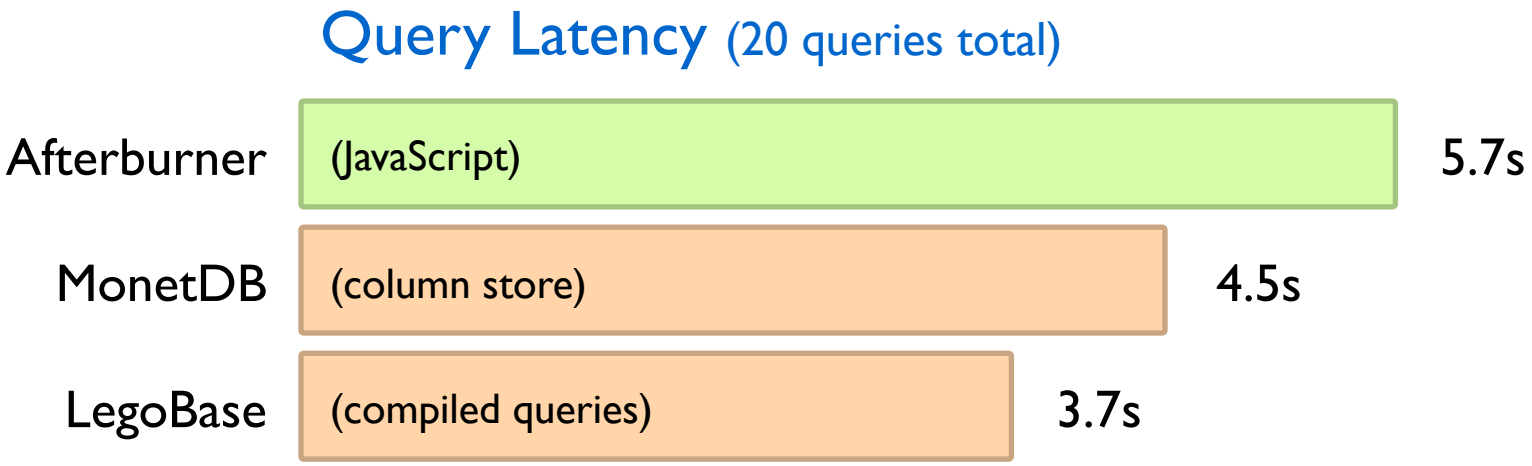

Note: we're comparing a browser tab vs. native OS execution!

Kareem El Gebaly and Jimmy Lin. In-Browser Interactive SQL Analytics with Afterburner. *SIGMOD 2017 Demo.*

Kareem El Gebaly and Jimmy Lin. In-Browser Split-Execution Support for Interactive Analytics in the Cloud. arXiv:1804.08822, 2018.

What would an in-browser JavaScript search engine look like?

What about a database?

What about neural networks?


TensorFlow.js

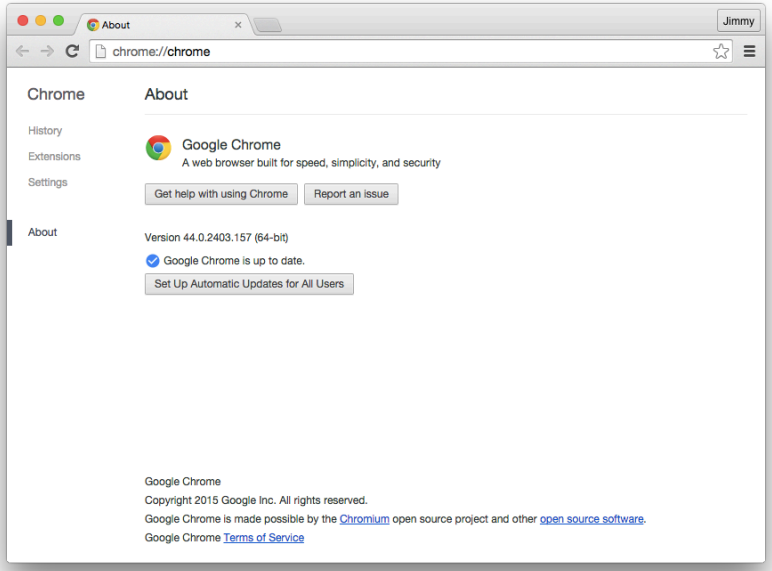

POSENET

Real-time Human Pose Estimation in the browser.

GO TO DEMO! CODE

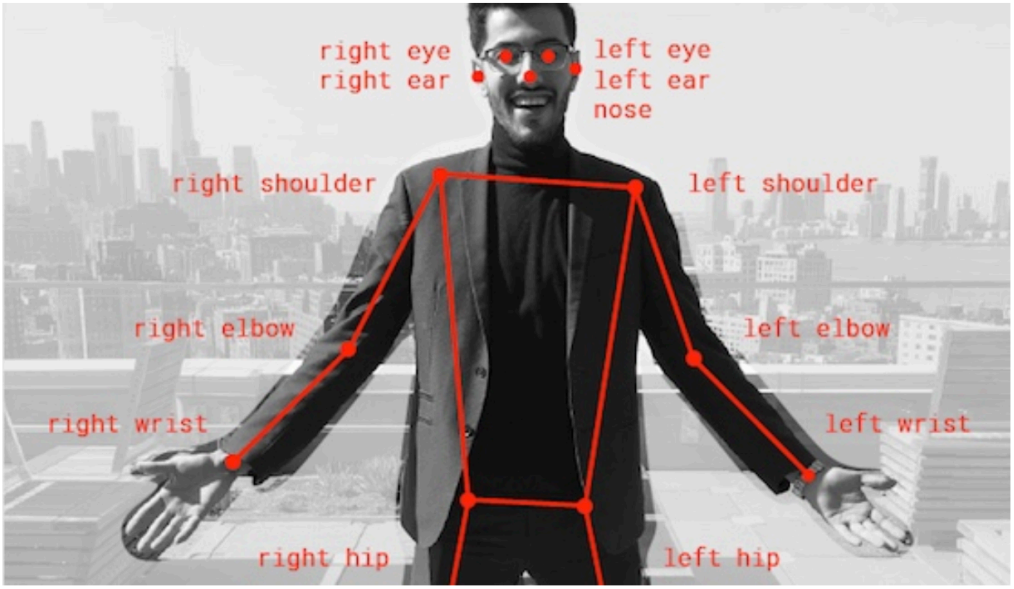

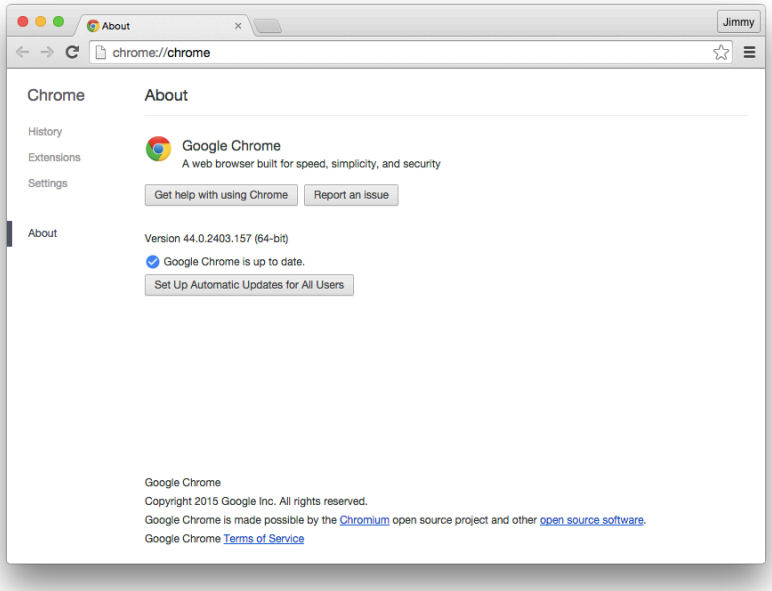

Vision applications – input directly from webcam…

Text applications – where do word embeddings come from?

IndexDB to the rescue!

(the same trick for storing postings lists)

How well does it work?

tl; dr – *surprisingly usable!*

Simple CNN for sentiment analysis

	Latency (ms) / batch			
	1	32	64	128
PyTorch				
Desktop GPU (Ubuntu 16.04)	2.9	3.0	3.1	3.1
Desktop CPU (Ubuntu 16.04)	4.3	43	86	130
Chrome Browser				
Desktop GPU (Ubuntu 16.04)	30	56	100	135
Desktop CPU (Ubuntu 16.04)	783	47900	110000	253000
MacBook Pro GPU (MacOS 10.13)	33	180	315	702
MacBook Pro CPU (MacOS 10.13)	779	56300	126000	297000
iPad Pro (iOS 11)	170	472	786	1283
Nexus 6P (Android 8.1.0)	103	541	1117	1722
iPhone 6 (iOS 11)	400	1336	3055	7324

Yiyun Liang, Zhucheng Tu, Laetitia Huang, and Jimmy Lin. CNNs for NLP in the Browser: Client-Side Deployment and Visualization Opportunities. *NAACL 2018 Demo.*

Technically interesting, perhaps… but:

Who cares?

The Balance of Computing

Where does computing happen?

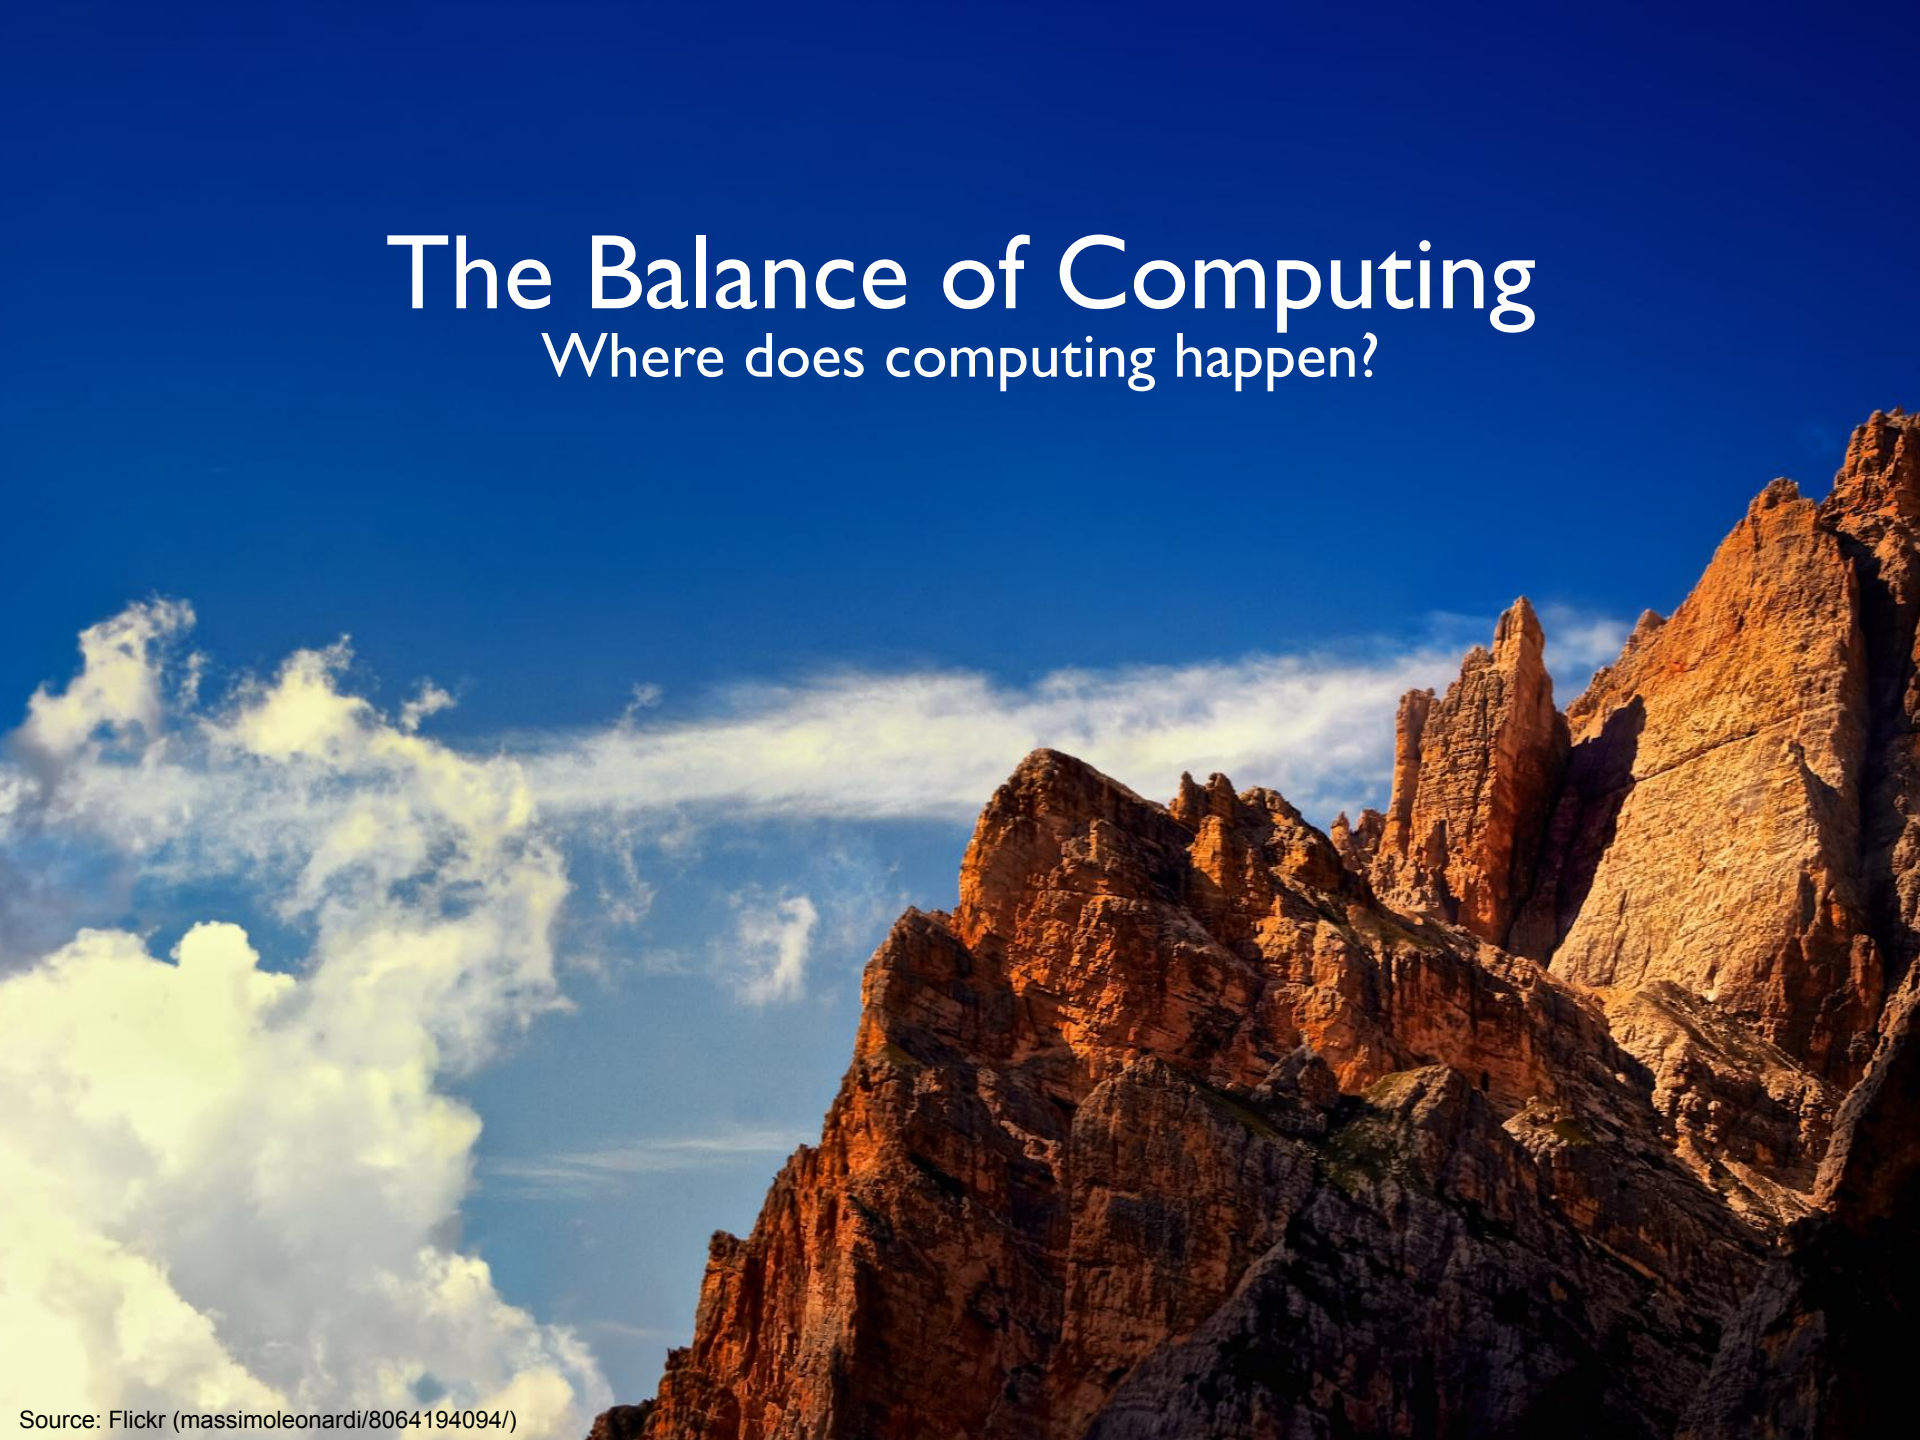

Source: Flickr (massimoleonardi/8064194094/)

Google bing

Baidu 百度 Яндекс

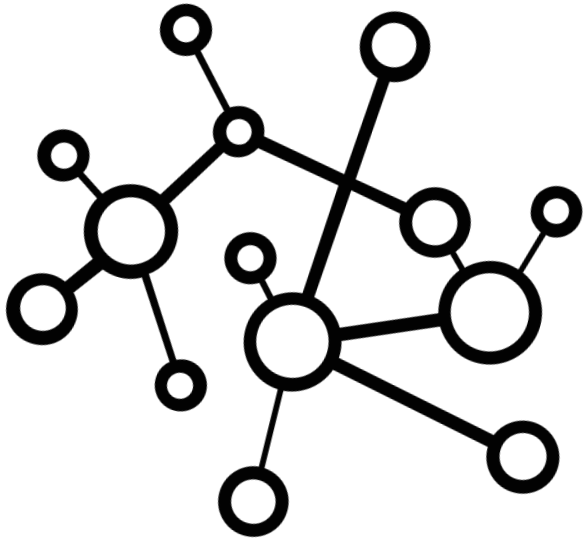

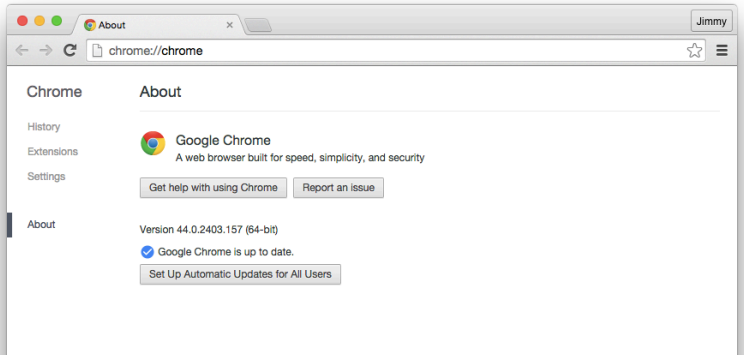

Google bing

Baidu 百度 Яндекс

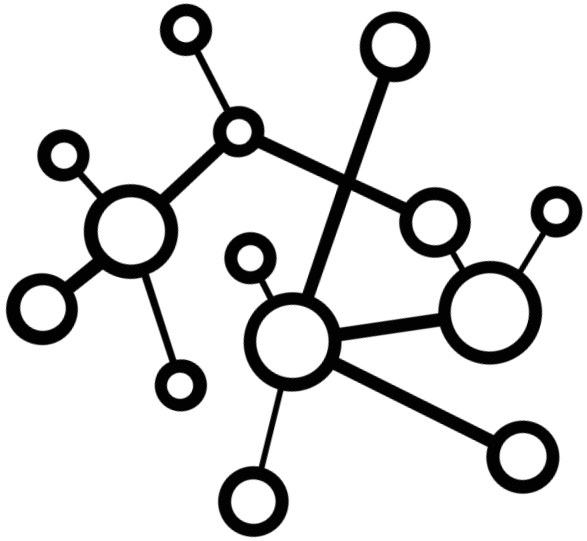

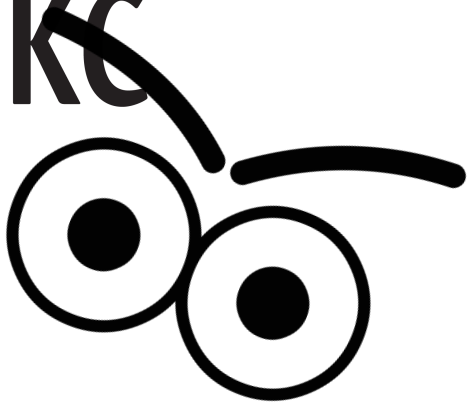

Male…

Geeky interests…

Traveling to Italy…

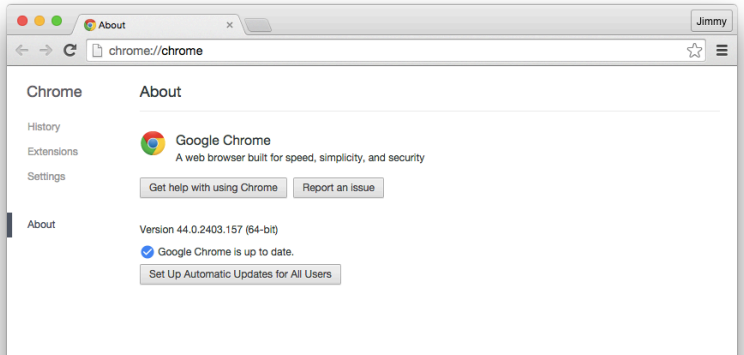

May have cancer…

What!?

Brief Report

FREE

March 2017

Evaluation of the Feasibility of Screening Patients for Early Signs of Lung Carcinoma in Web Search Logs

Ryen W. White, PhD[1]; Eric Horvitz, MD, PhD[1]

» Author Affiliations | Article Information

JAMA Oncol. 2017;3(3):398-401. doi:10.1001/jamaoncol.2016.4911

Key Points

Question Are statistical m
carcinoma in advance of a

Findings In this modeling
later input queries that pr
ods can help identify peop
identify new risk factors (e
activity and geographic lo

Meaning Pattern analysis
risk factors and for framin

Abstract

Importance A statistical m
diagnosis via analysis of la
across the United States.

Objective To evaluate the feasibility of screening patients at risk of lung carcinoma via analysis of signals from online search activity.

Design, Setting, and Participants We identified people who issue special queries that provide strong evidence of a recent diagnosis of lung carcinoma. We then considered patterns of symptoms expressed as searches about concerning symptoms over several months prior to the appearance of the landmark web queries. We built statistical classifiers that predict the future appearance of landmark queries based on the search log signals. This was a retrospective log analysis of the online activity of millions of web searchers seeking health-related information online. Of web searchers who queried for symptoms related to lung carcinoma, some (n=5443 of 4 813 985) later issued queries that provide strong evidence of recent clinical diagnosis of lung carcinoma and are regarded as positive cases

Results The statistical classifier predicting the future appearance of landmark web queries based on search log signals identified searchers who later input queries consistent with a lung carcinoma diagnosis, with a true-positive rate ranging from 3% to 57% for false-positive rates ranging from 0.00001 to 0.001, respectively. The methods can be used to identify people at highest risk up to a year in advance of the inferred diagnosis time. The 5 factors associated with the highest relative risk (RR) were evidence of family history (RR = 7.548; 95% CI, 3.937-14.470), age (RR = 3.558; 95% CI, 3.357-3.772), radon (RR = 2.529; 95% CI, 1.137-5.624), primary location (RR = 2.463; 95% CI, 1.364-4.446), and occupation (RR = 1.969; 95% CI, 1.143-3.391). Evidence of smoking (RR = 1.646; 95% CI, 1.032-2.260) was important but not top-ranked, which was due to the difficulty of identifying smoking history from search terms.

Google bing

Baidu 百度 Яндекс

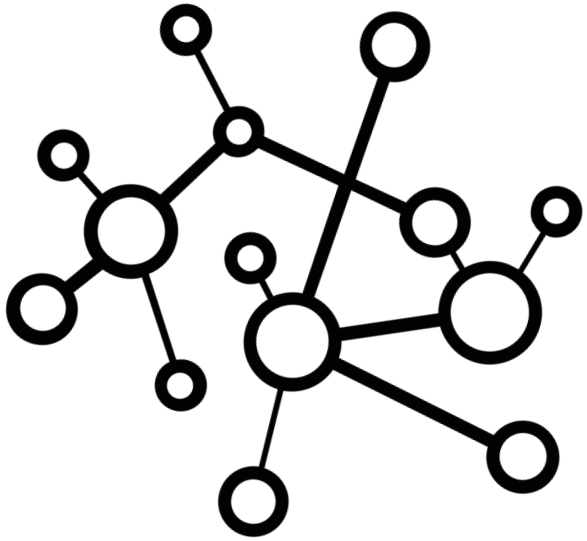

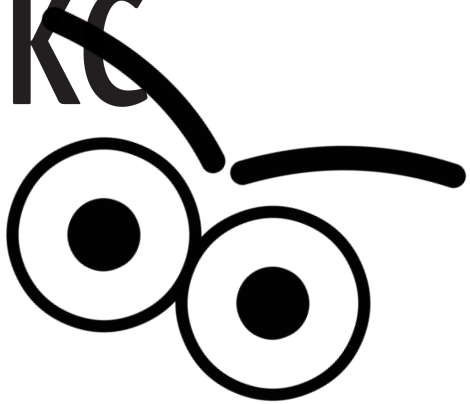

Male…
Geeky interests…
Traveling to Italy…

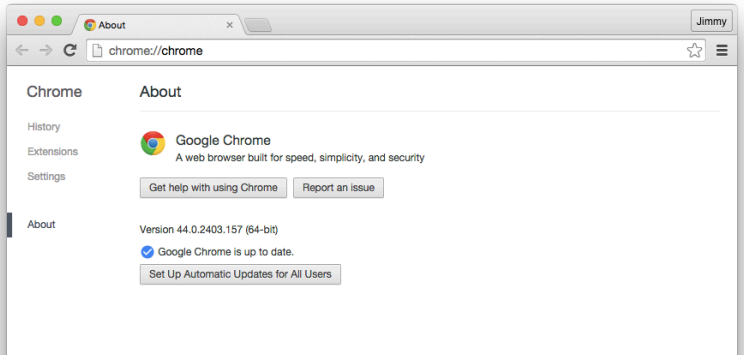

May have cancer…
Potential bankruptcy…
Risk of self-harm…

100s of billions of web pages

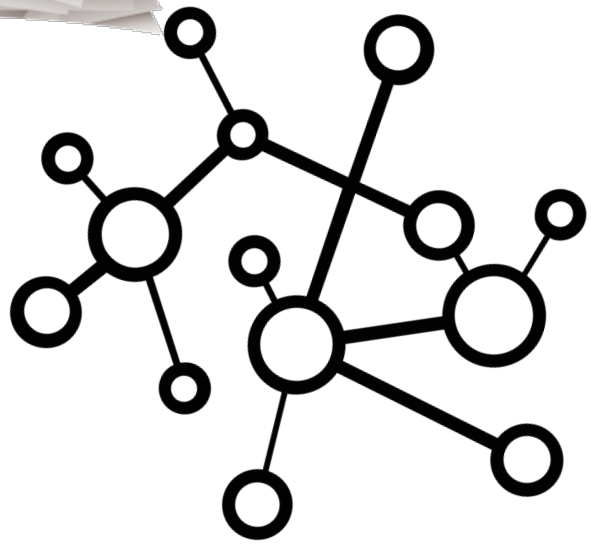

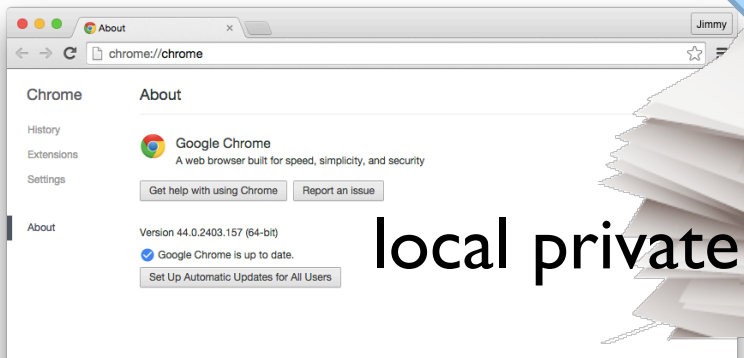

Small focused collection
(e.g., health information)

local private search

Local Private Search

Information Interactions

Let's store a local copy of Wikipedia! #

When was the last time you searched Google and ended up at Wikipedia?

When was the last time you used a search engine to find a page you recently visited? *

While we're at it, let's cache (a portion of) the web!

Let's archive every web page you've ever visited!

Feel free to divulge your inner-most secrets to a neural network!

Other advantages?

Jimmy Lin. The Sum of All Human Knowledge in Your Pocket: Full-Text Searchable Wikipedia on a Raspberry Pi. *JCDL 2015.*

* Sarah K. Tyler and Jaime Teevan. Large Scale Query Log Analysis of Re-Finding. *WSDM 2010.*

Google bing

Baidu 百度 Яндекс

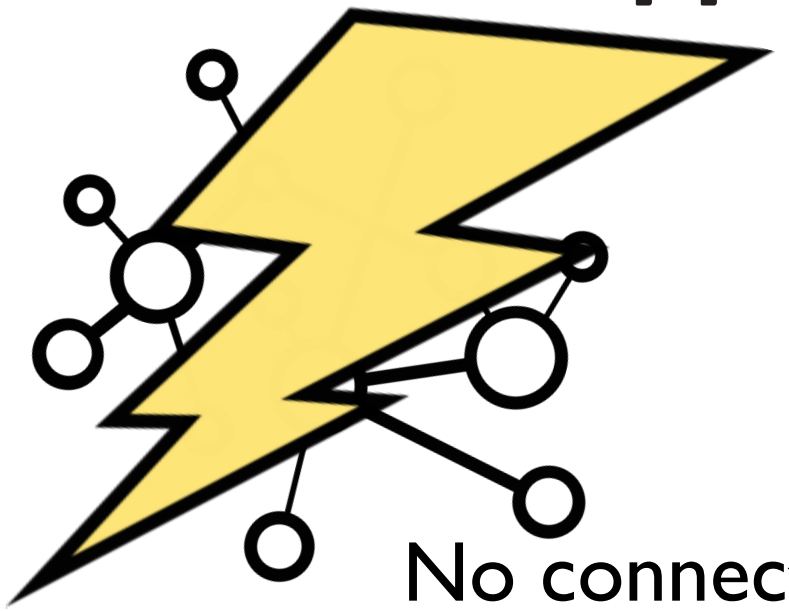

No connectivity?

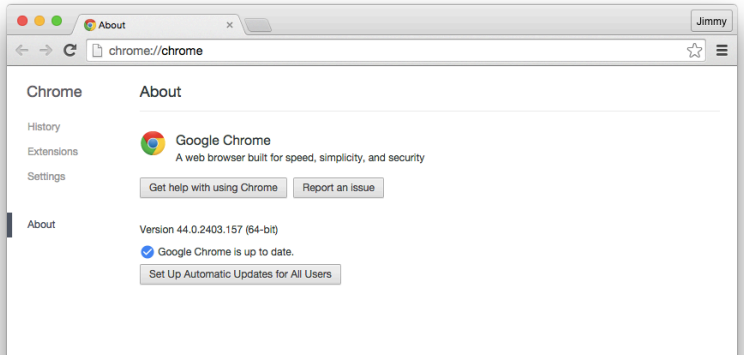

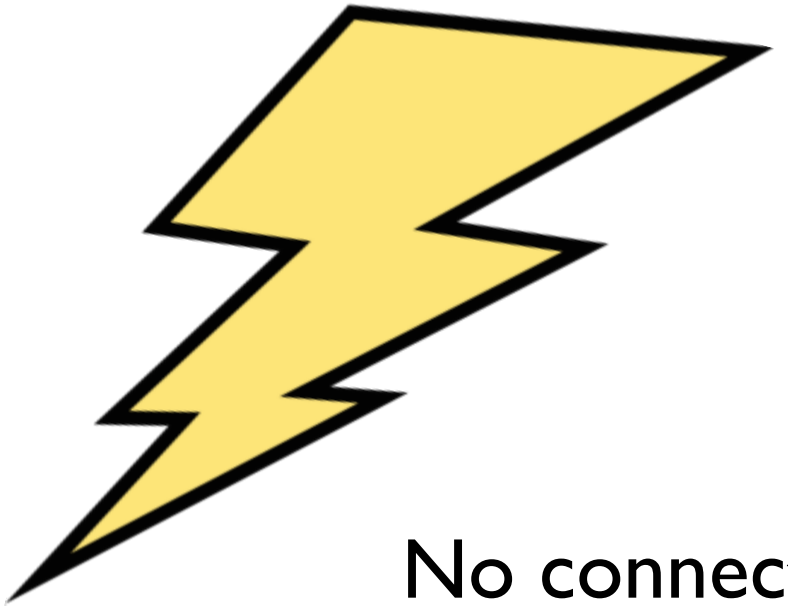

No connectivity?

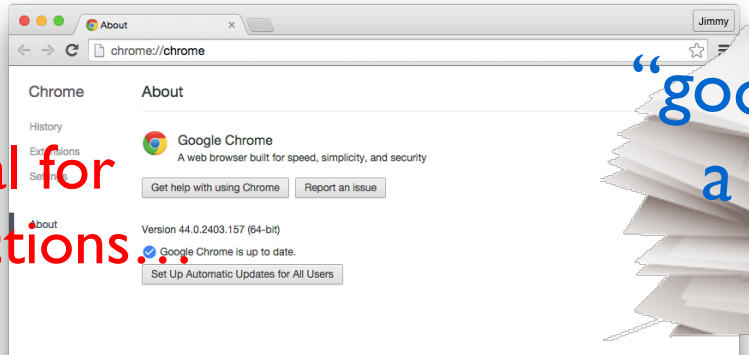

"good thing I kept a local copy!"

Additional potential for low-latency interactions

But wait… none of this actually *requires* JavaScript!

Yes, but JavaScript makes it incredibly convenient!

"It's just a webpage" deployment!

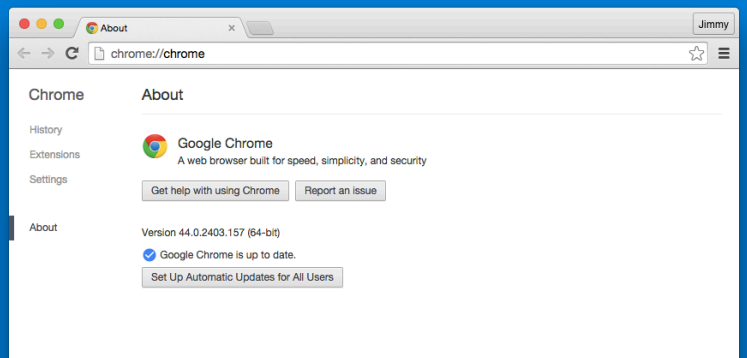

100s of billions of web pages

general search

private search

Interesting {split, cooperative} execution architectures?

Small focused collection

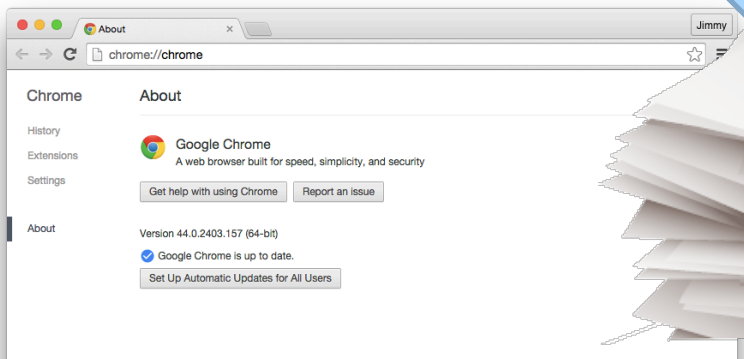

Split Execution for Neural Networks!

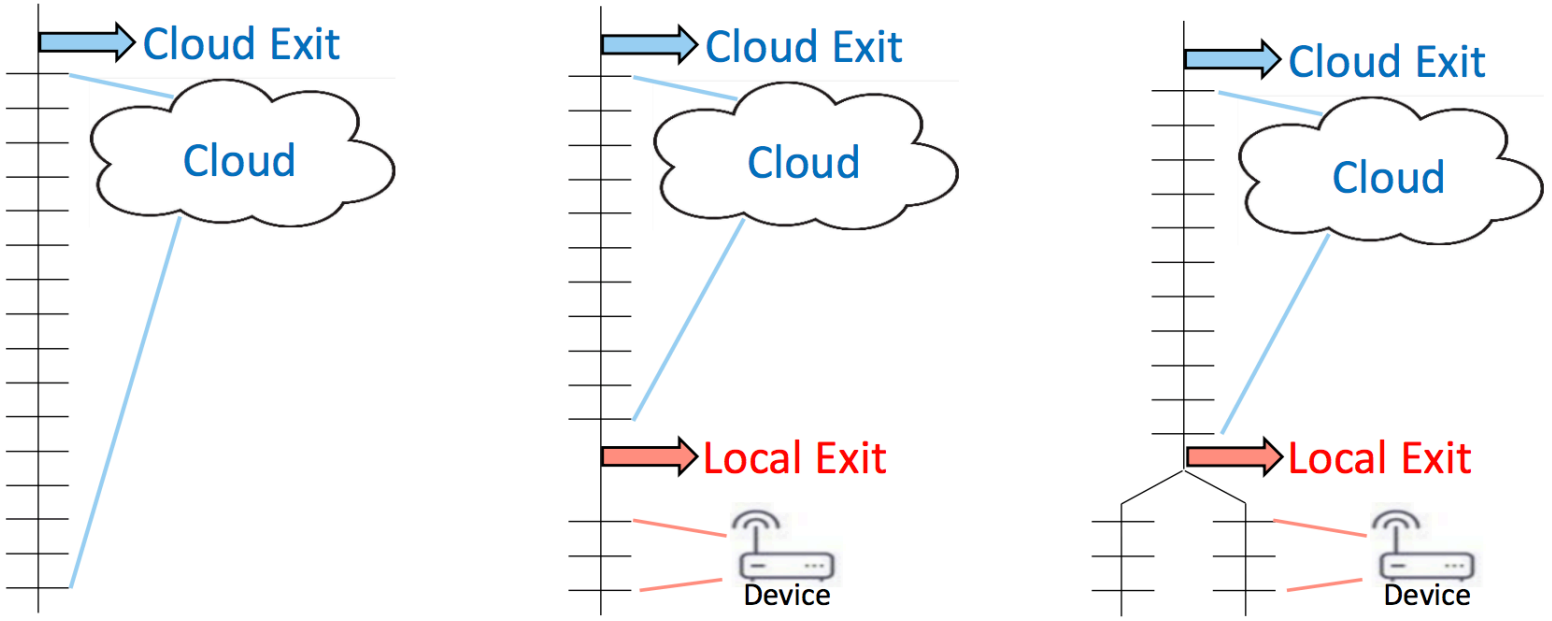

Surat Teerapittayanon, Bradley McDanel, and H.T. Kung. Distributed Deep Neural Networks over the Cloud, the Edge and End Devices. *ICDCS 2017.*

Split Execution for Neural Networks!

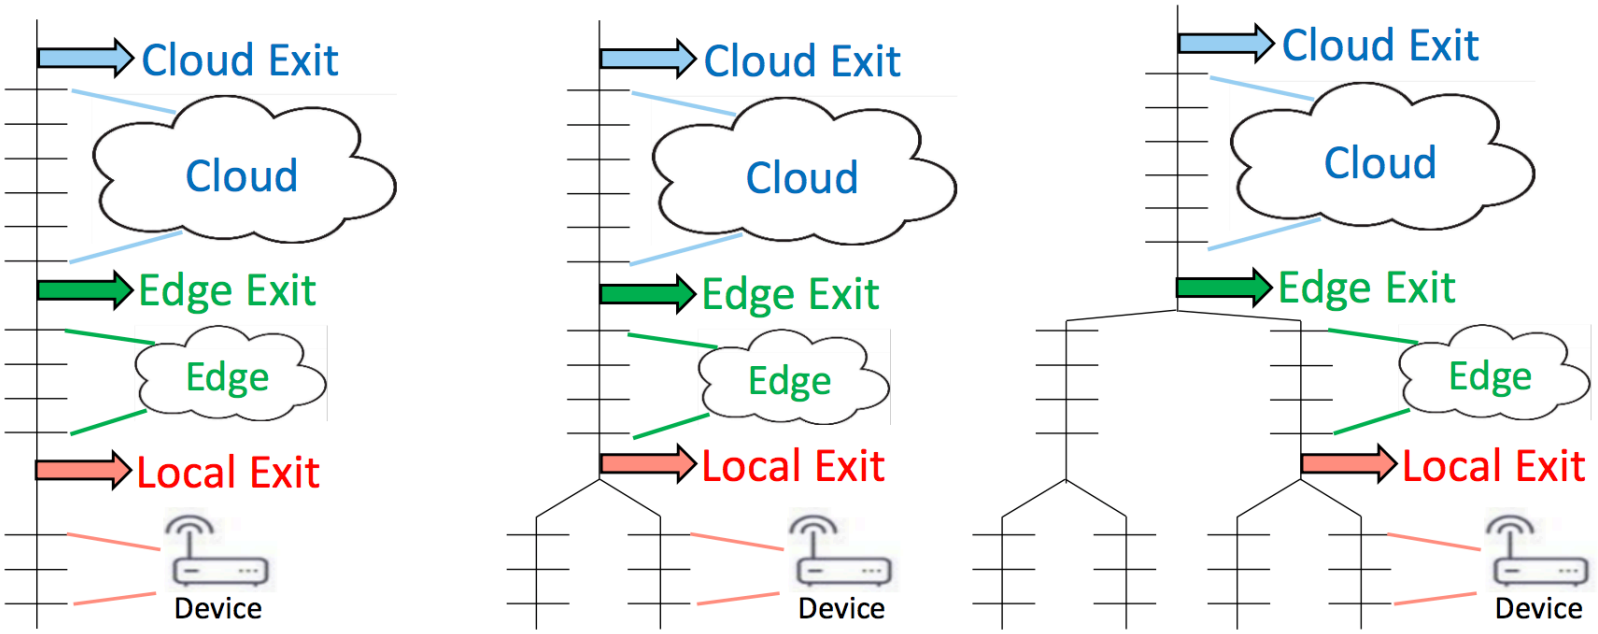

Hmm… does this remind you of something?

Surat Teerapittayanon, Bradley McDanel, and H.T. Kung. Distributed Deep Neural Networks over the Cloud, the Edge and End Devices. *ICDCS 2017.*

The Balance of Computing

Where does computing happen?

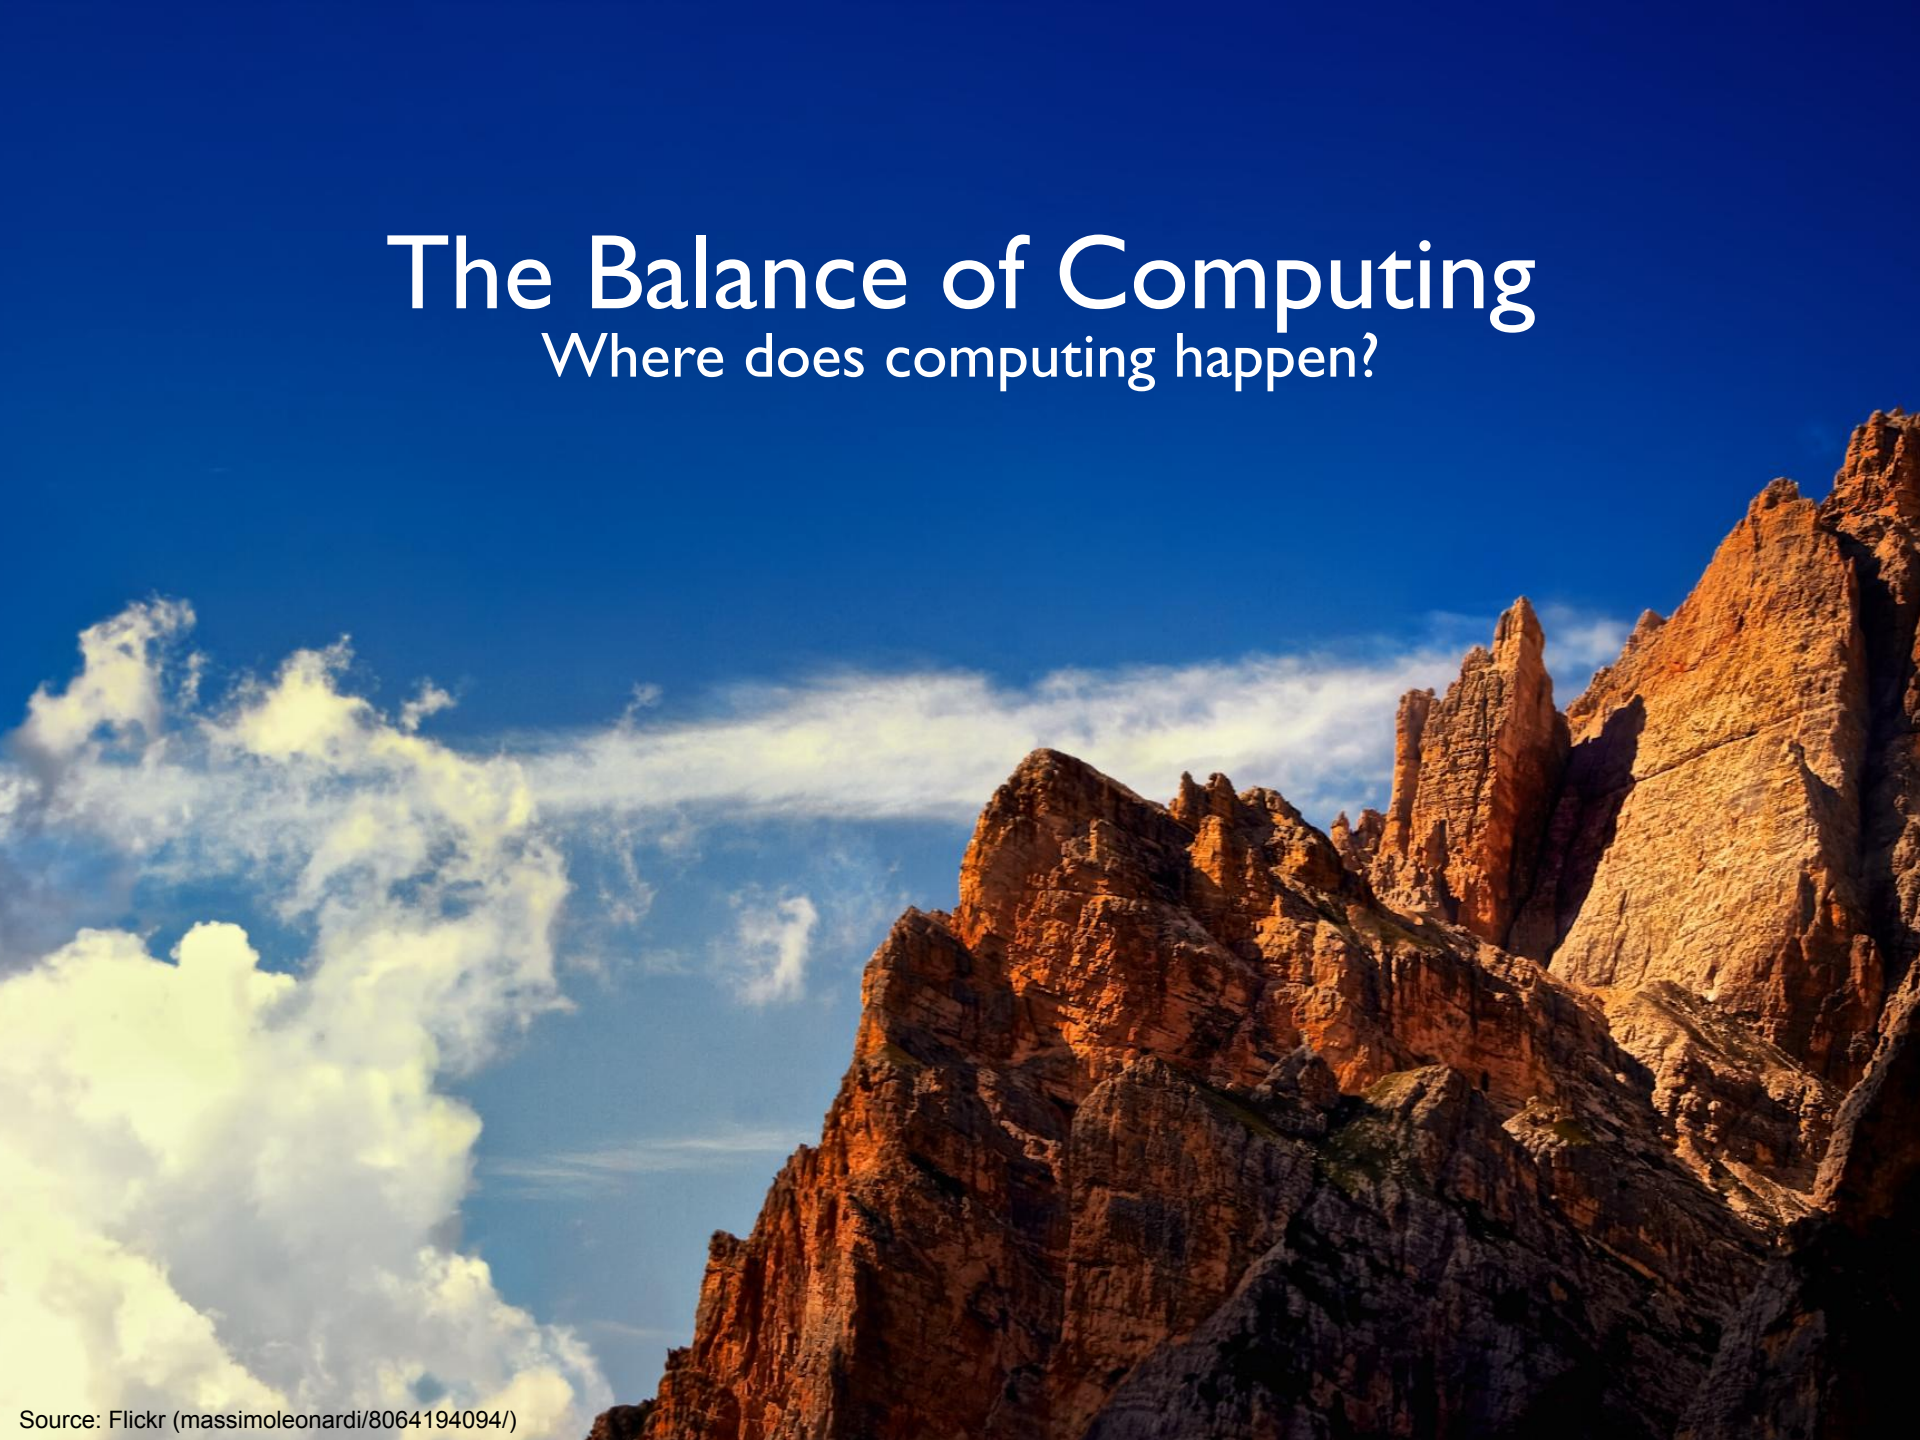

Source: Flickr (massimoleonardi/8064194094/)

Source: Flickr (massimoleonardi/8064194094/)

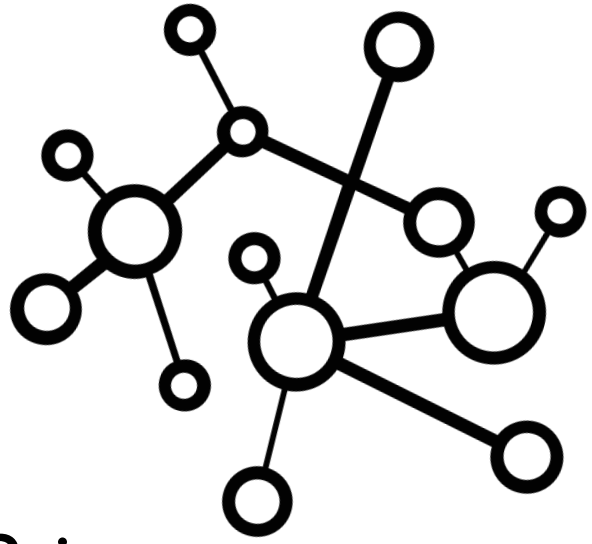

Prizm

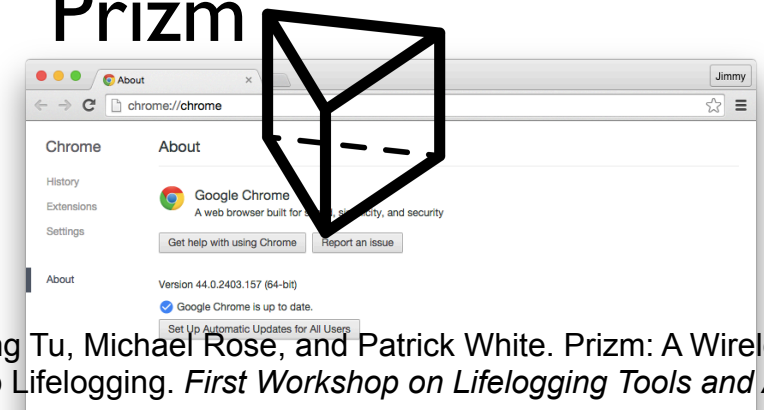

Jimmy Lin, Zhucheng Tu, Michael Rose, and Patrick White. Prizm: A Wireless Access Point for Proxy-Based Web Lifelogging. *First Workshop on Lifelogging Tools and Applications, 2016.*

What is Prizm?

Wireless access point deployed on a Raspberry Pi

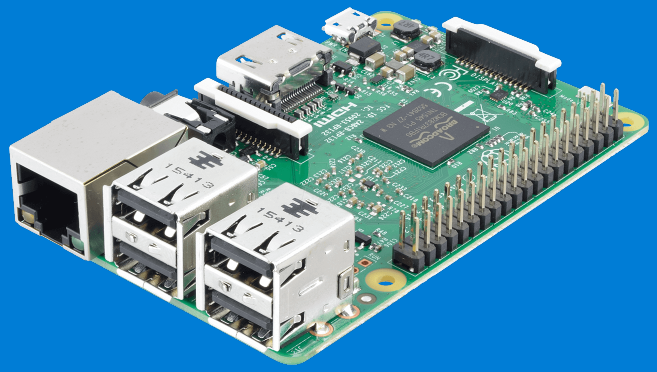

Proxies all http(s) requests and captures all activity

Original application was lifelogging

Jimmy Lin, Zhucheng Tu, Michael Rose, and Patrick White. Prizm: A Wireless Access Point for Proxy-Based Web Lifelogging. *First Workshop on Lifelogging Tools and Applications, 2016.*

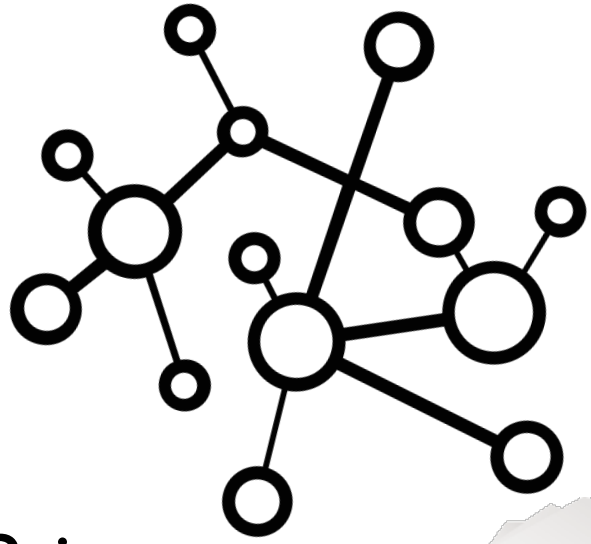

Prizm

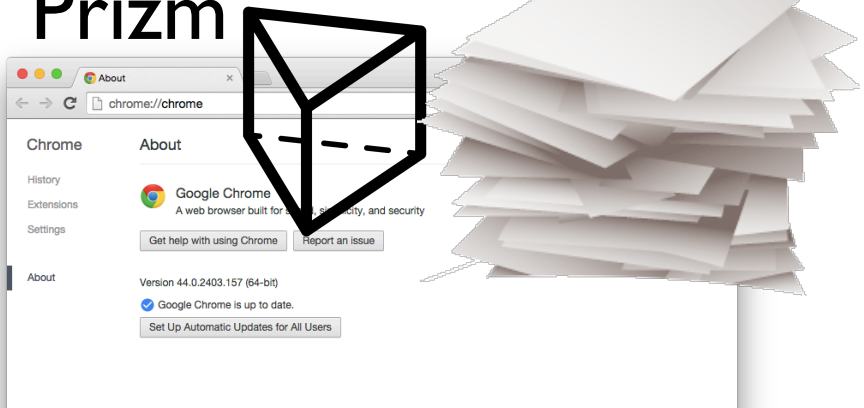

Transparently build local collection by observing user behavior…

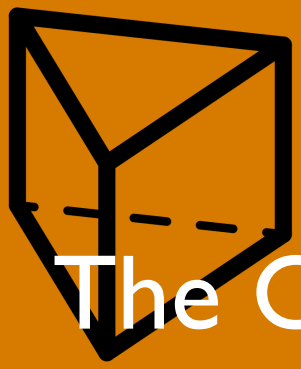

The Complete Vision

Prizm serves as the user's intelligent agent that controls all network connections with the outside world

Can this request be handled locally?
Does this interaction leak sensitive information?
…

Brokerage of private log data (with user consent)

Source: https://qz.com/1013634/spacex-launch-watch-live-as-elon-musks-rocket-company-launches-bulgarias-first-communications-satellite/

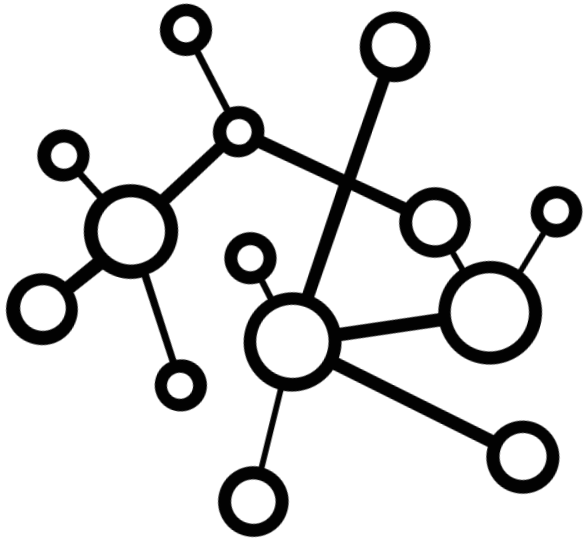

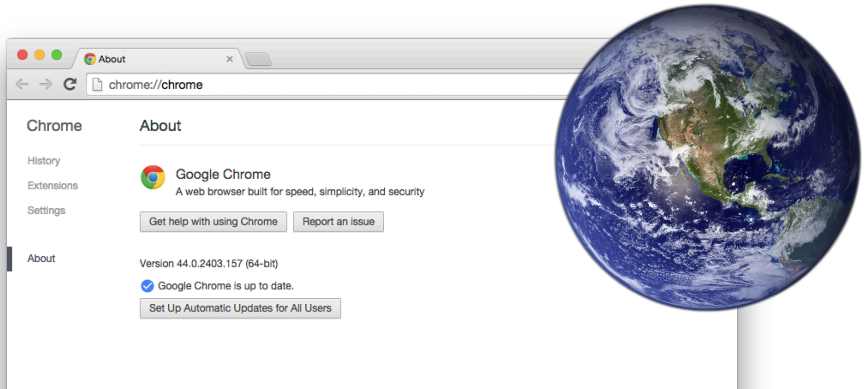
About
chrome://chrome
Chrome
History
Extensions
Settings
About
About
Google Chrome
A web browser built for speed, simplicity, and security
Get help with using Chrome
Report an issue
Version 44.0.2403.157 (64-bit)
Google Chrome is up to date.
Set Up Automatic Updates for All Users

Source: Flickr (massimoleonardi/8064194094/)

Source: NASA/JPL

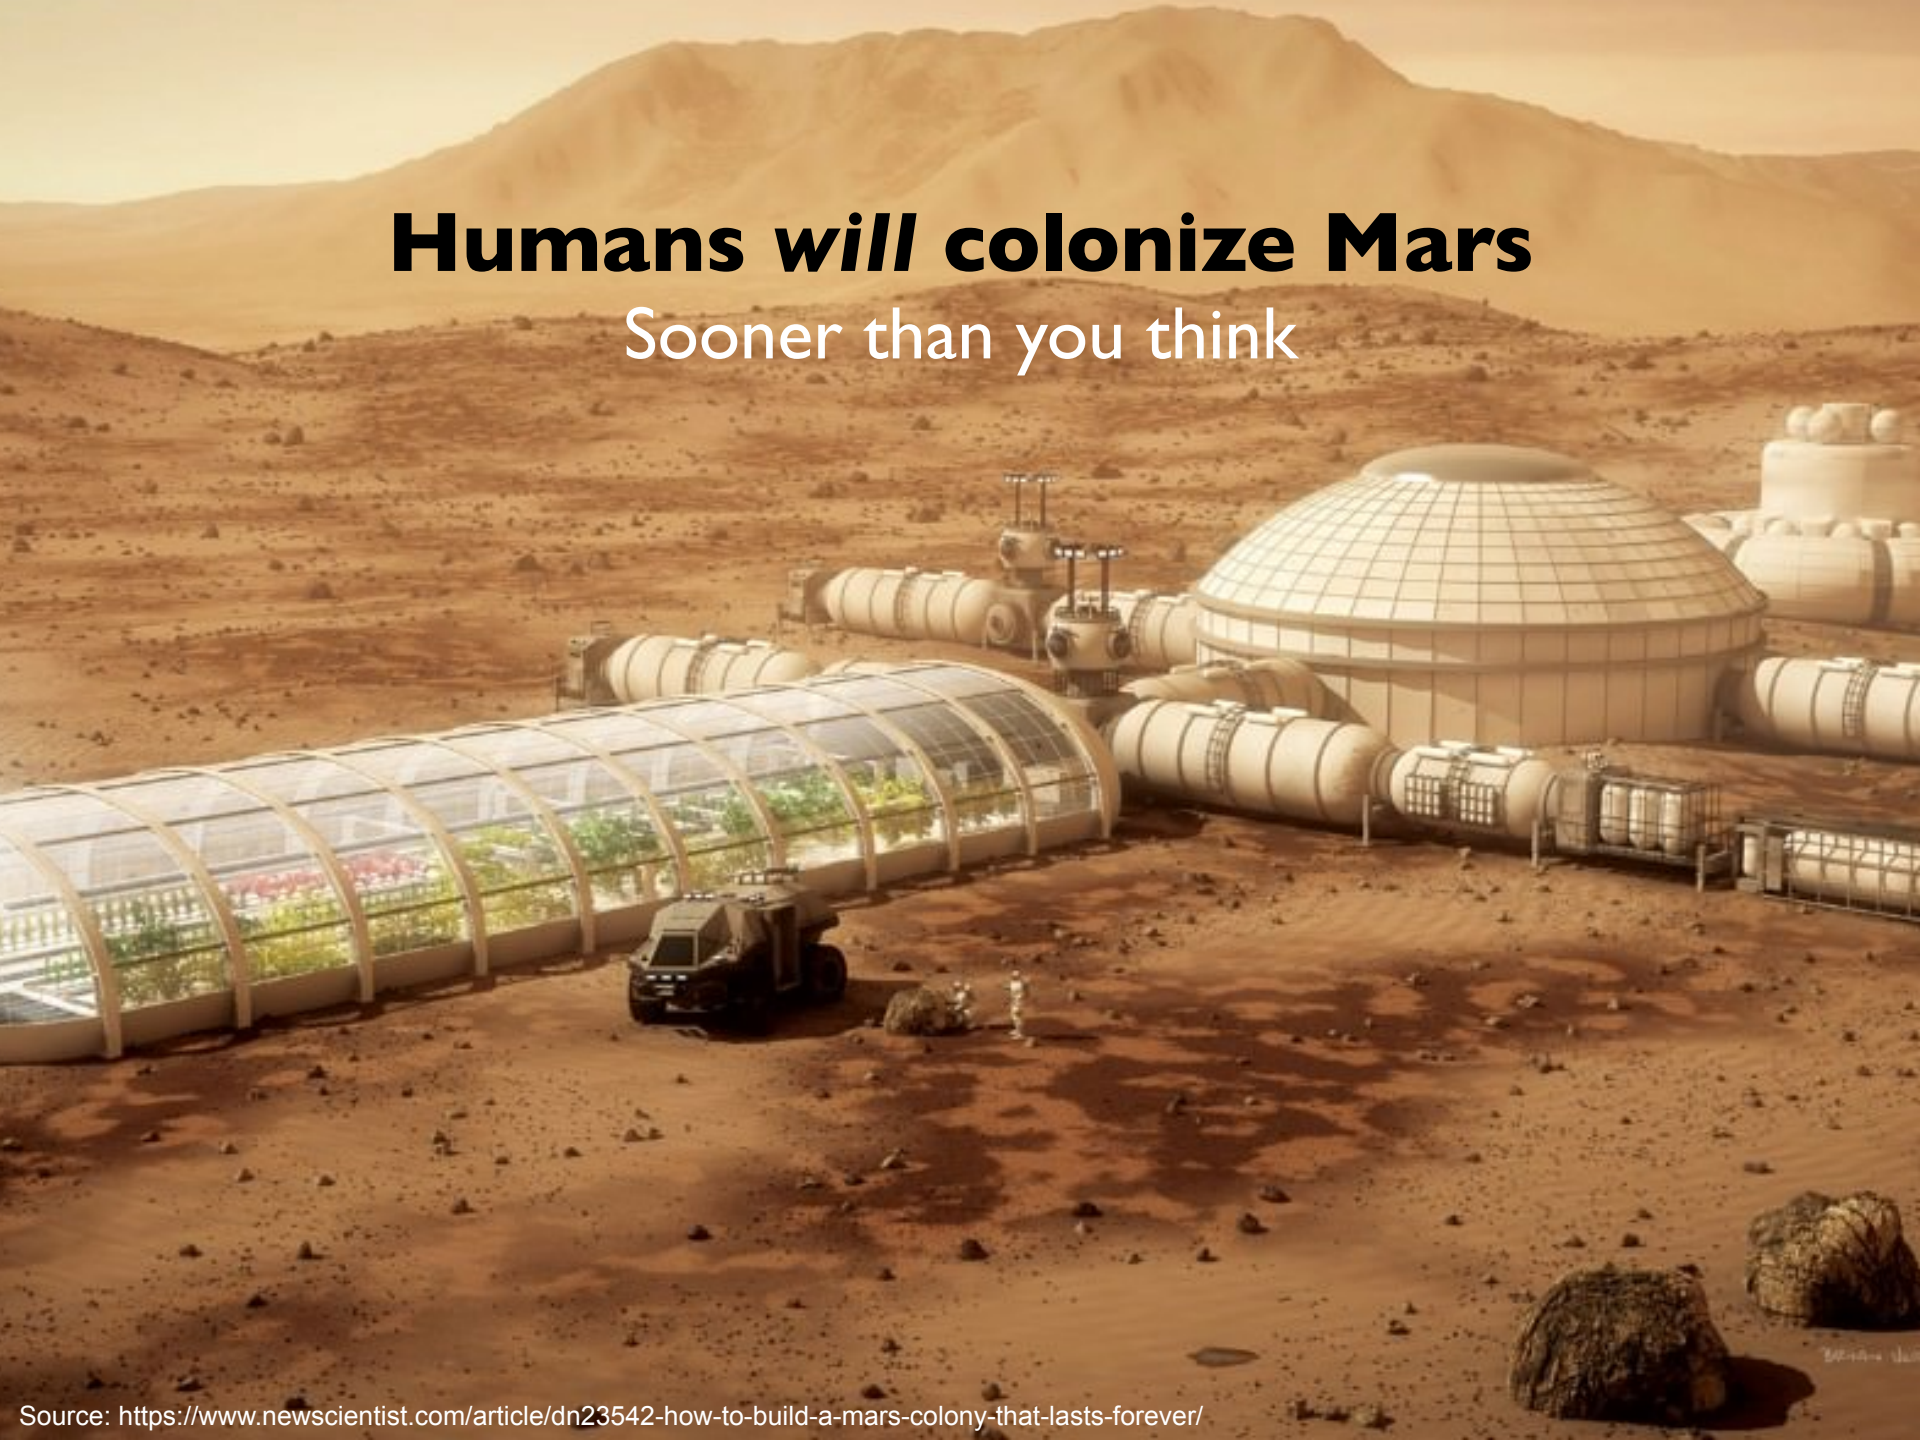

Source: https://www.newscientist.com/article/dn23542-how-to-build-a-mars-colony-that-lasts-forever/

Elon Musk Charts Path to Colonizing Mars Within a Decade

By Robin Seemangal • 06/06/16 9:10am

SpaceX
@SpaceX

Planning to send Dragon to M
2018. Red Dragons will infor
architecture, details to come

RETWEETS 15,988 LIKES 26,640

11:50 AM - 27 Apr 2016

16K 27K

Buzz Aldrin developing a 'master plan' to colonize Mars within 25 years

Aldrin and the Florida Institute of Technology are pushing for a Mars settlement by 2039, the 70th anniversary of his own Apollo 11 moon landing

Florida Tech's president, Anthony J Catanese, left, talks with Apollo 11 astronaut Buzz Aldrin as he shows him the campus on Thursday in Melbourne, Florida Photograph: Craig Rubadoux/AP

Source: https://www.theguardian.com/science/2015/aug/27/buzz-aldrin-colonize-mars-within-25-years
Source: https://twitter.com/SpaceX/status/725351354537906176
Source: http://observer.com/2016/06/elon-musk-charts-path-to-colonizing-mars-within-a-decade/

“Mars can’t just be a one-shot mission” – Buzz Aldrin

“The Pilgrims on the Mayflower came here to live and stay. They didn’t wait around Plymouth Rock for the return trip, and neither will people building up a population and a settlement [on Mars].”

Source: *Mayflower in Plymouth Harbor* by William Halsall (1882)

Needs

Produce breathable air
Grow food
Build shelter
Mine fuel and materials

"Staying alive"

Conduct science

Connect with family and friends
Engage in leisure activities
Search the web

"Staying sane"

Searching the web should be as easy from Mars as it is from Marseille!

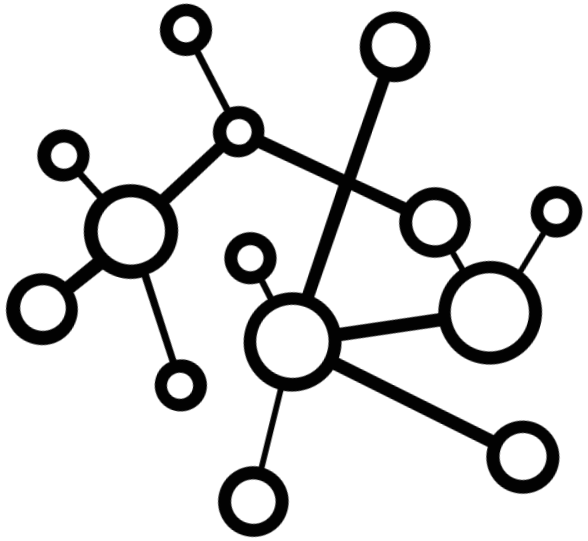

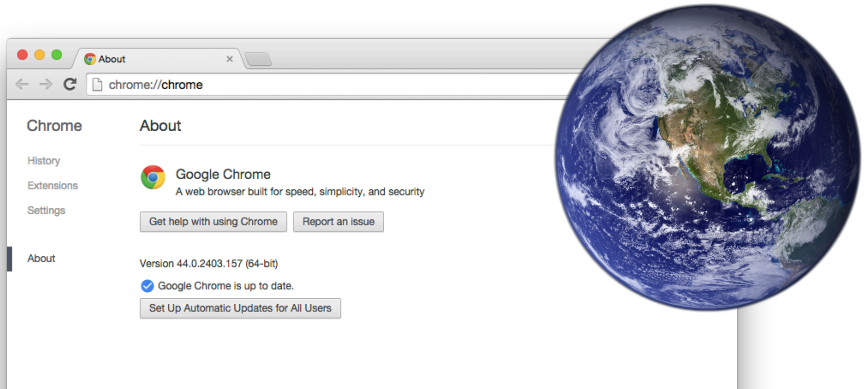
About
chrome://chrome
Chrome
History
Extensions
Settings
About
About
Google Chrome
A web browser built for speed, simplicity, and security
Get help with using Chrome
Report an issue
Version 44.0.2403.157 (64-bit)
Google Chrome is up to date.
Set Up Automatic Updates for All Users

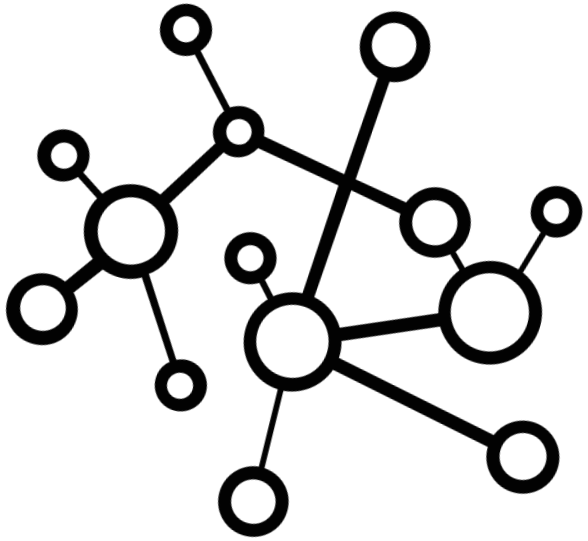

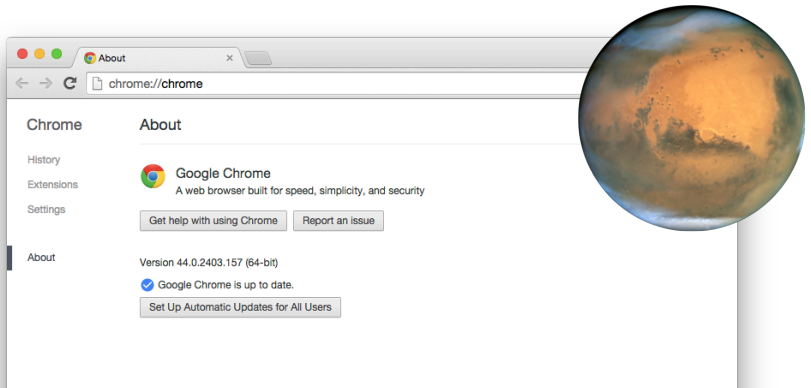
About
chrome://chrome
Chrome
History
Extensions
Settings
About
About
Google Chrome
A web browser built for speed, simplicity, and security
Get help with using Chrome
Report an issue
Version 44.0.2403.157 (64-bit)
Google Chrome is up to date.
Set Up Automatic Updates for All Users

The fundamental problem: Latency

speed of light: 2-24 minutes
rockets: 5-10 months

Bandwidth is "reasonable"

Lunar Laser Communications Demonstration:
622-Mbps downlink, 20-Mbps uplink

SneakerNet on rockets: Easily PBs

This is exactly how we start!

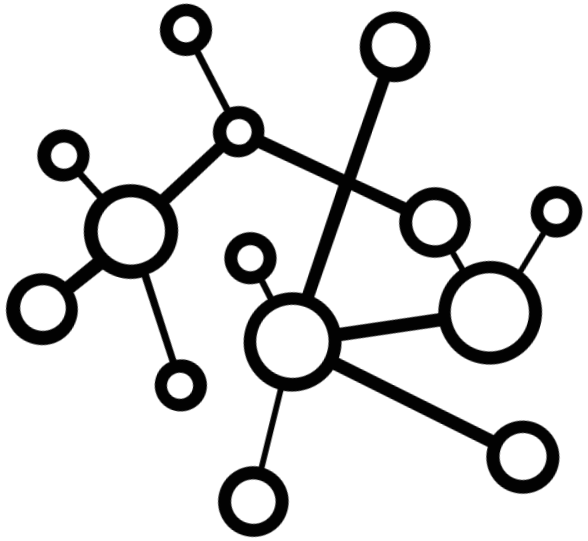

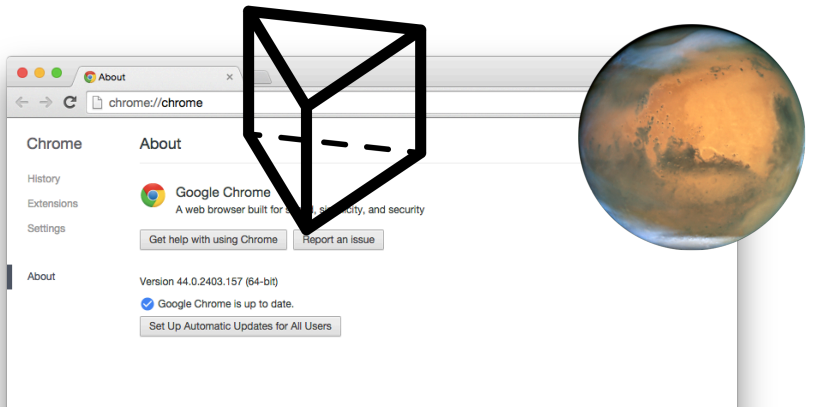
About
chrome://chrome
Chrome
About
History
Extensions
Settings
About
Google Chrome
Get help with using Chrome
Report an issue
Version 44.0.2403.157 (64-bit)
Google Chrome is up to date.
Set Up Automatic Updates for All Users

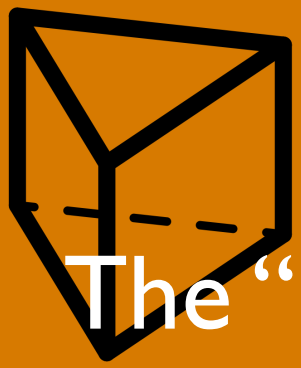

The “Martian Prizm”

Manages all inbound and outbound network traffic

Pipelines user actions over to Earth to hide latencies

Predictively pre-fetches data to populate the Martian cache

Naive

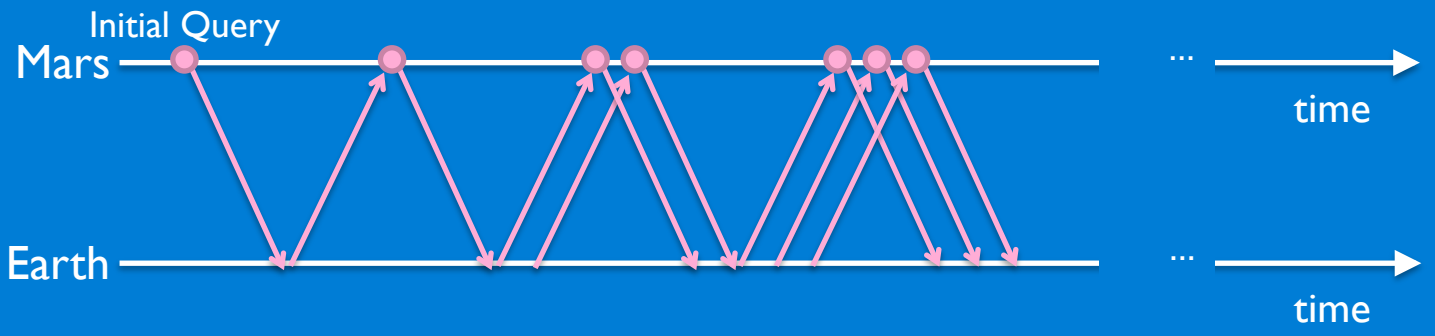

Mars cache + pipelining interactions

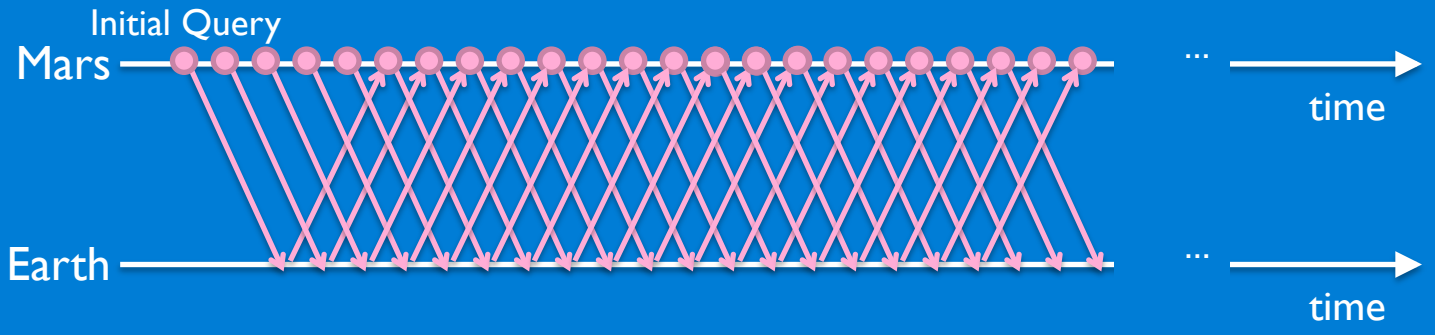

Charles L. A. Clarke, Gordon V. Cormack, Jimmy Lin, and Adam Roegiest. Ten Blue Links on Mars. *WWW 2017.*

Search from Mars is definitely feasible!

Search from Pluto, on the other hand…

For the truly skeptical…

Search from Mars ~
Search from regions on Earth with poor connectivity

Easter Island
Canadian Arctic
Villages in rural India

More "down to Earth" applications!

The Balance of Computing

Where does computing happen?

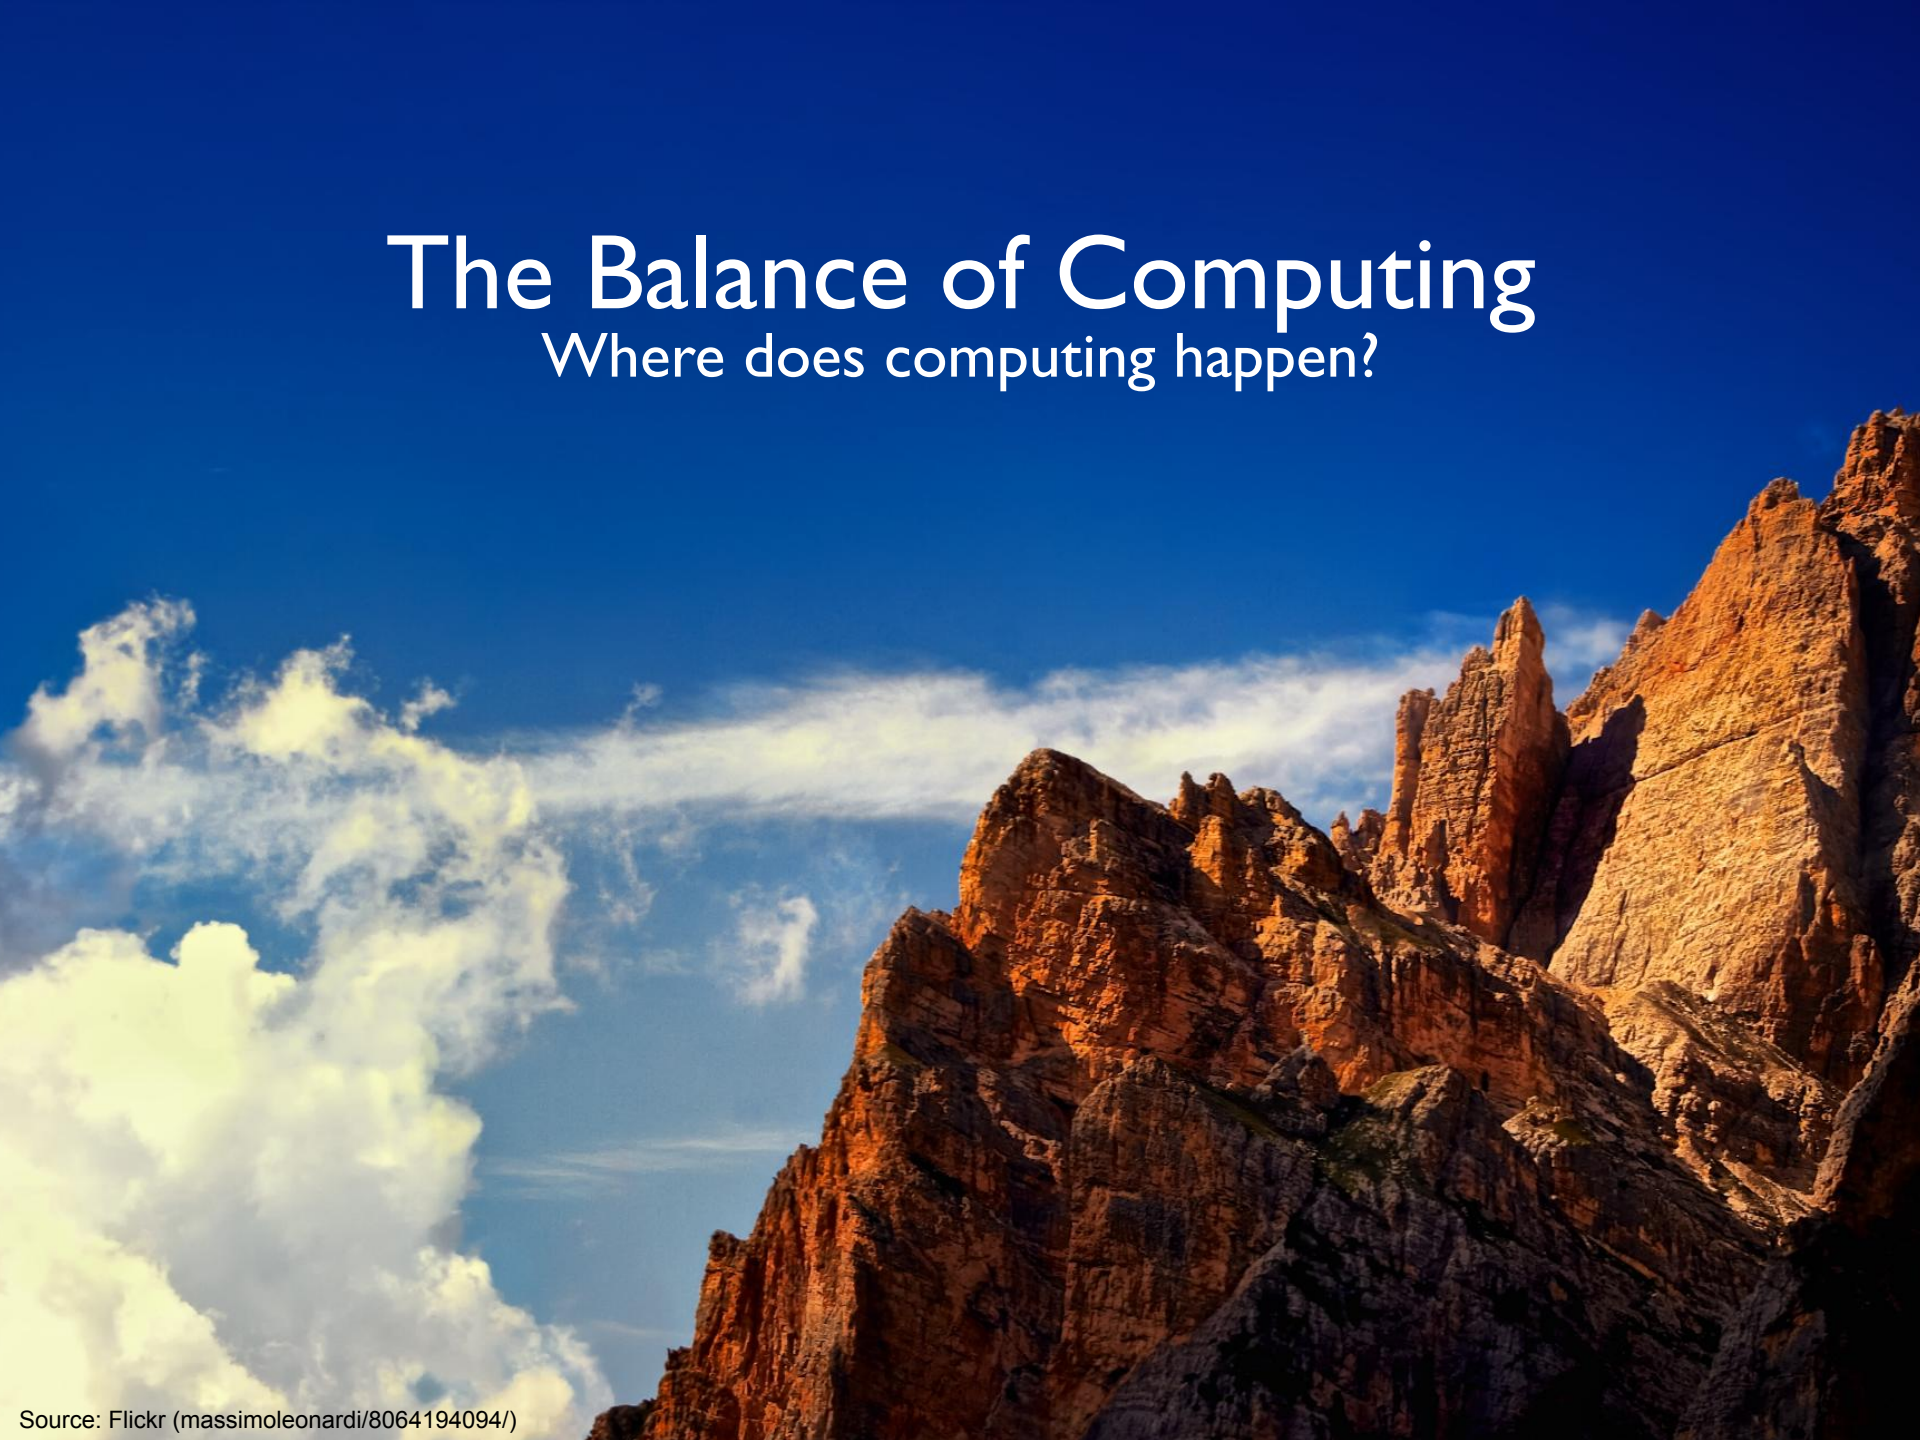

Source: Flickr (massimoleonardi/8064194094/)

Source: Flickr (michael_hendrickx/5327363552)

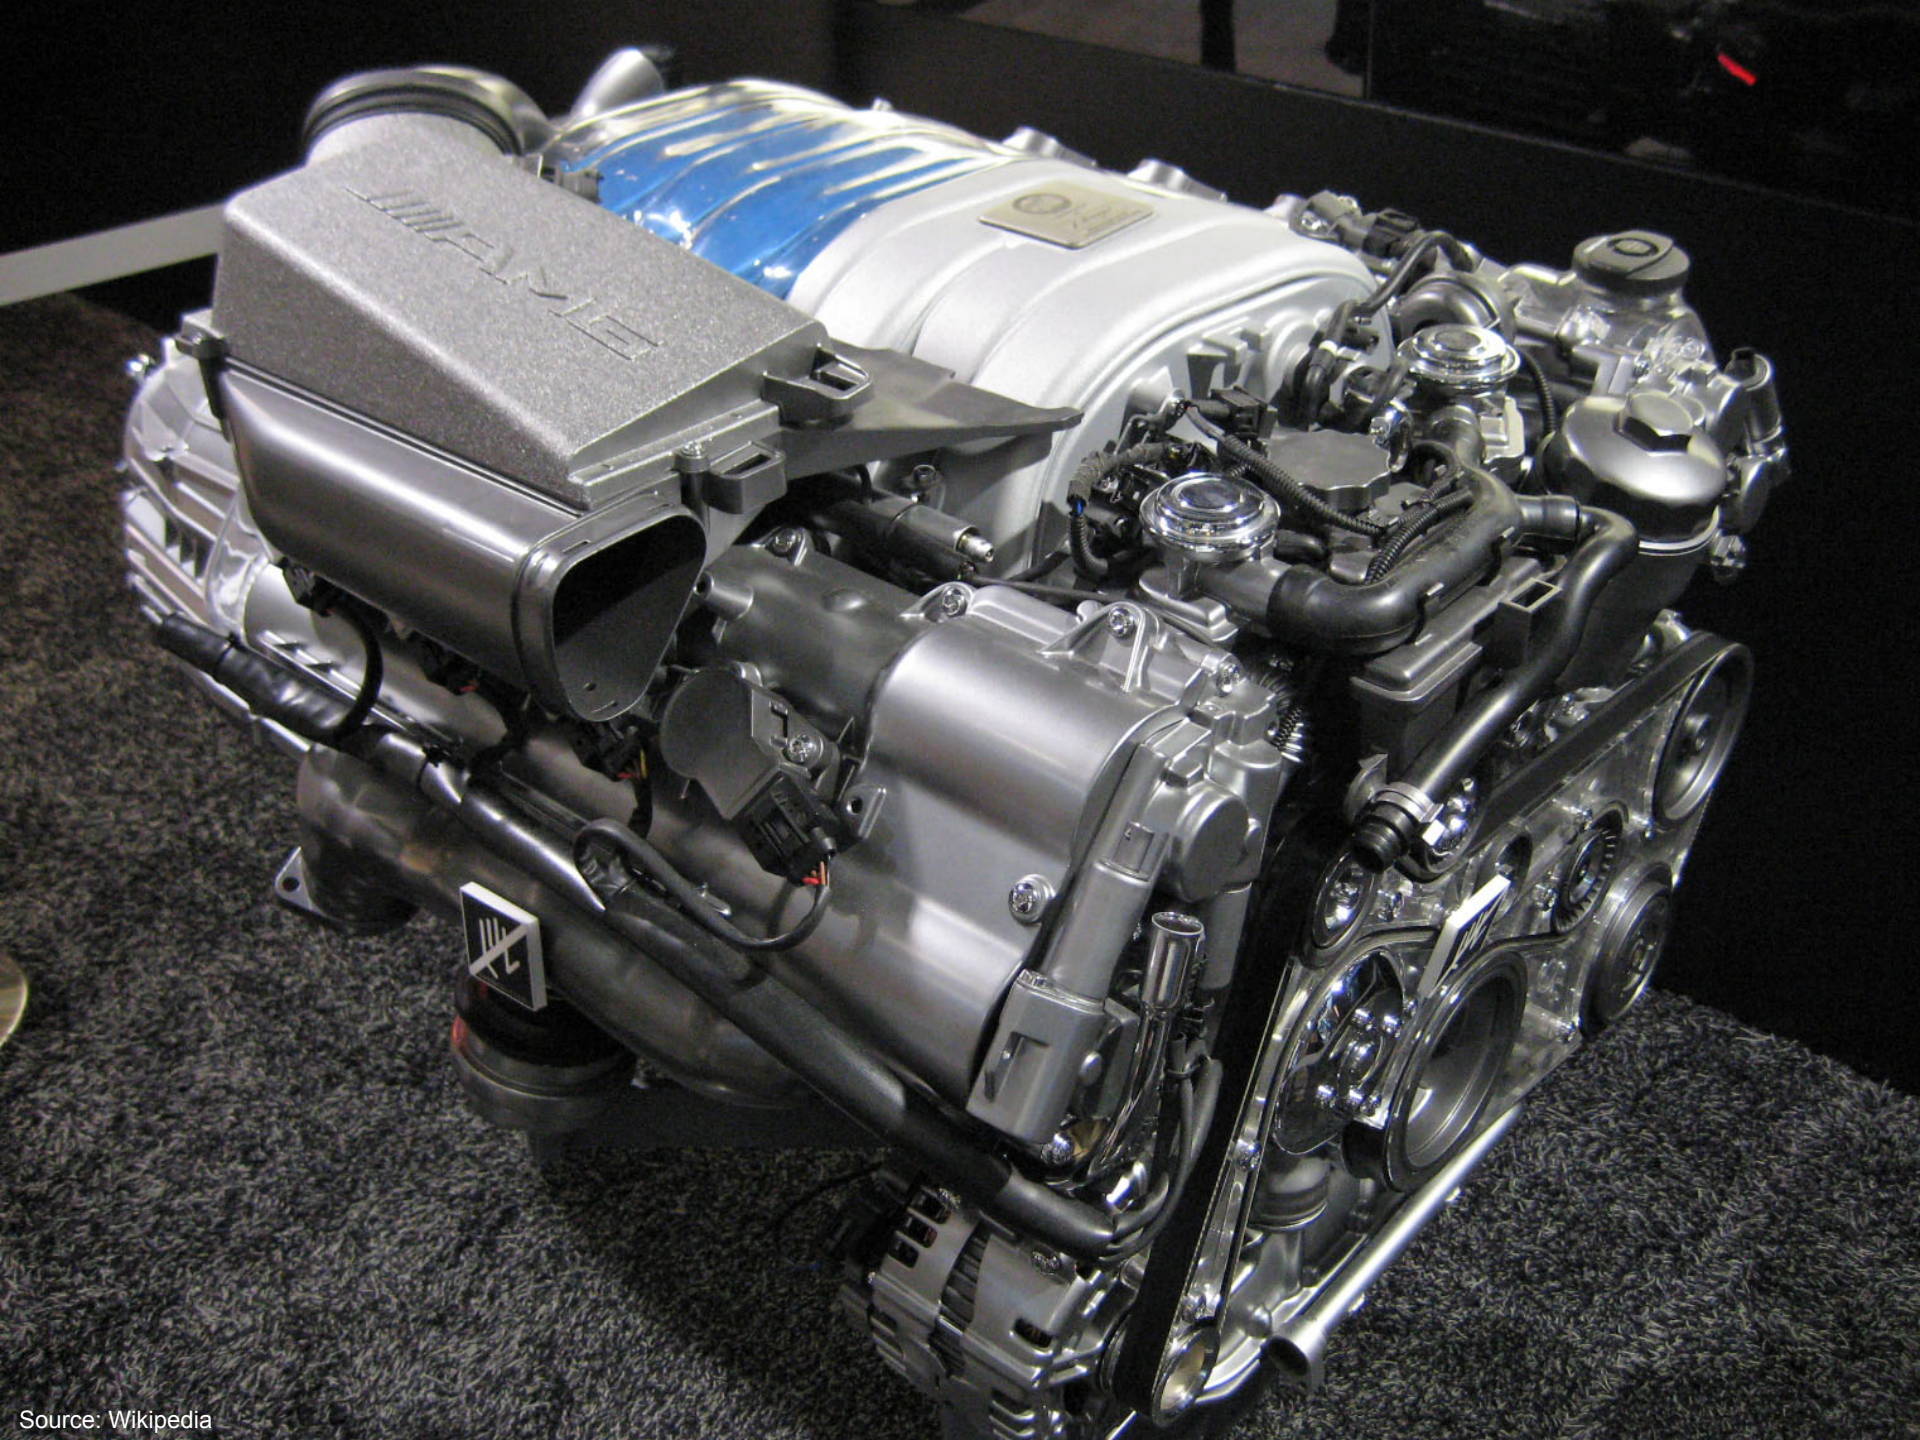

Source: Wikipedia

Source: https://qz.com/1013634/spacex-launch-watch-live-as-elon-musks-rocket-company-launches-bulgarias-first-communications-satellite/

On Batsh*t Crazy Ideas…

It is the customary fate of new truths to begin as heresies and to end as superstitions.

Thomas Henry Huxley. The Coming of Age of the Origin of Species. 1880.

They laughed at Columbus, they laughed at Fulton, they laughed at the Wright Brothers.
But they also laughed at Bozo the Clown.
Being laughed at does not mean you are right.

Carl Sagan. Broca's Brain: Reflections on the Romance of Science. Random House, 1979.

Shallit, Jeffrey (2005) Science, Pseudoscience, and The Three Stages of Truth.

On Batsh*t Crazy Ideas…

It is the customary fate of new truths to begin as heresies and to end as superstitions.

Thomas Henry Huxley. The Coming of Age of the Origin of Species. 1880.

They laughed at Columbus, they laughed at Fulton, they laughed at the Wright Broth
But they also laughed at Bozo the
Being laughed at does not mean yo

Carl Sagan. Broca's Brain: Reflections on the Romance of Scien

Shallit, Jeffrey (2005) Science, Pseudoscience, and The Three Stages of Truth.

Questions?

Source: Flickr (massimoleonardi/8064194094/)