Learning in a Distributed and Heterogeneous Environment

Martin Jaggi
EPFL Machine Learning and Optimization Laboratory
mlo.epfl.ch

Inria - EPFL joint Workshop - Paris - Feb 15th

Machine Learning Methods to Analyze Large-Scale Data

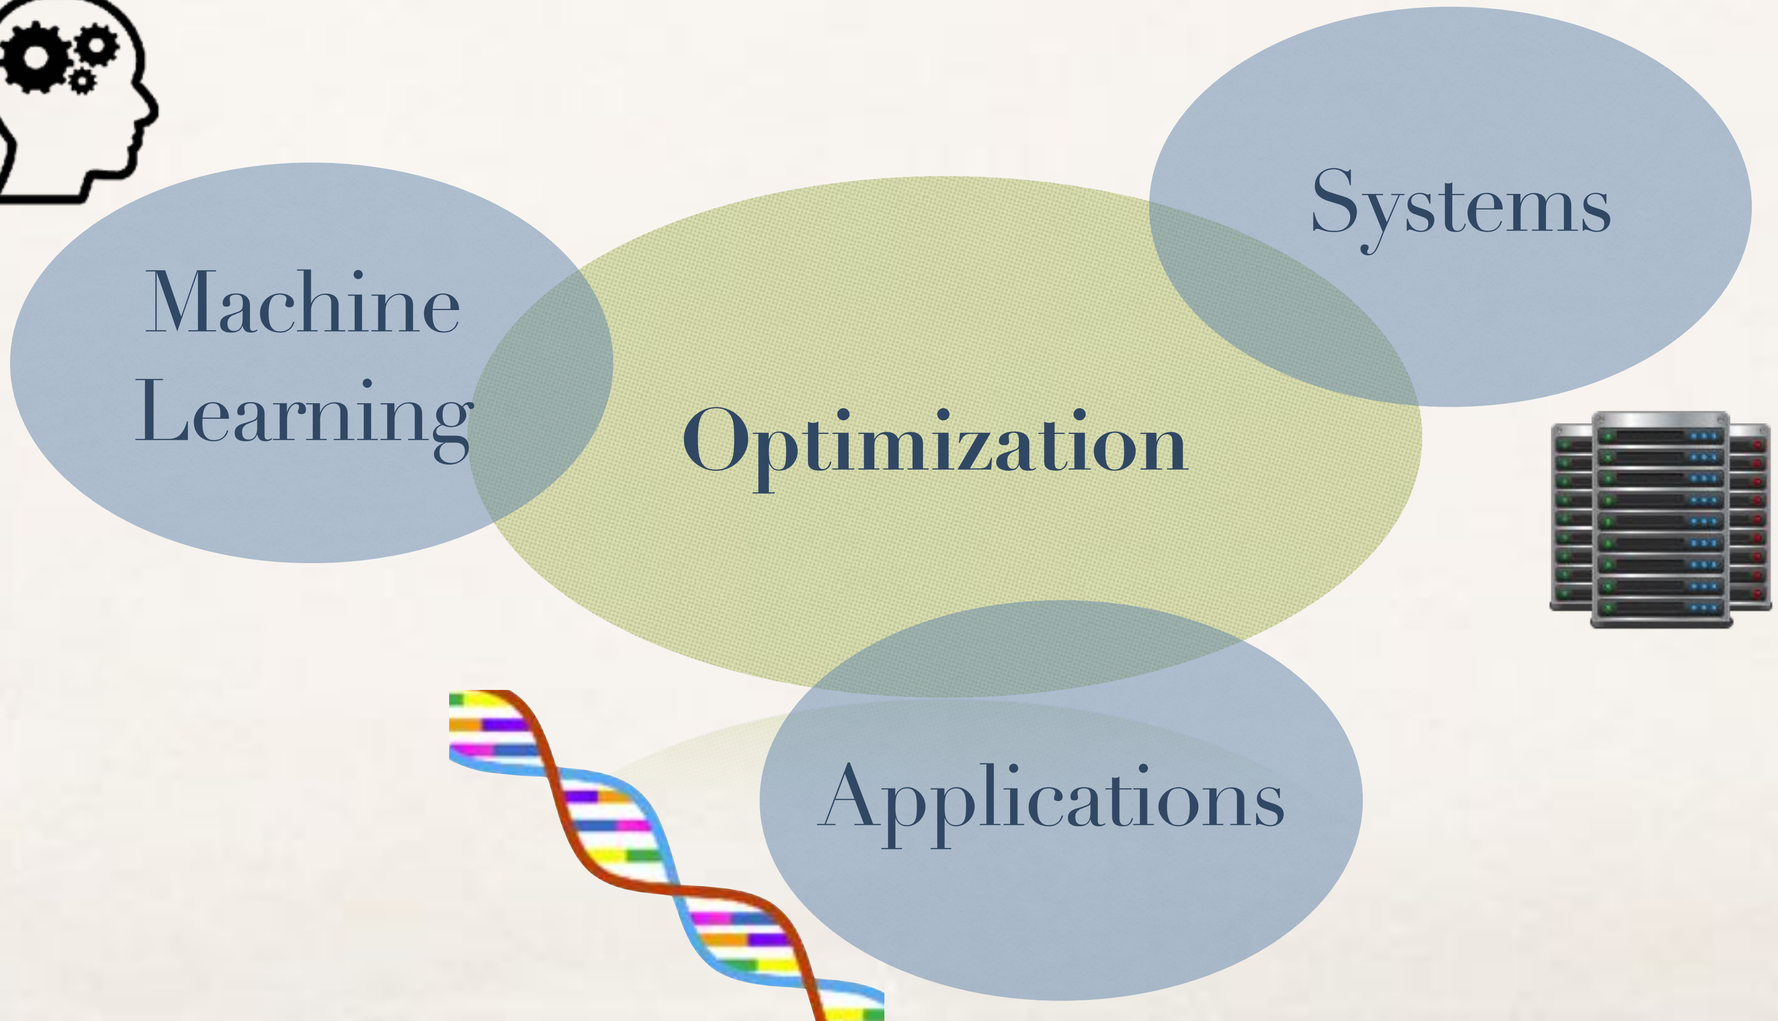

Machine Learning Systems

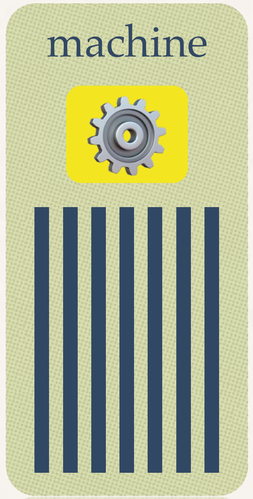

Foto: Florian Hirzinger

Machine Learning Systems

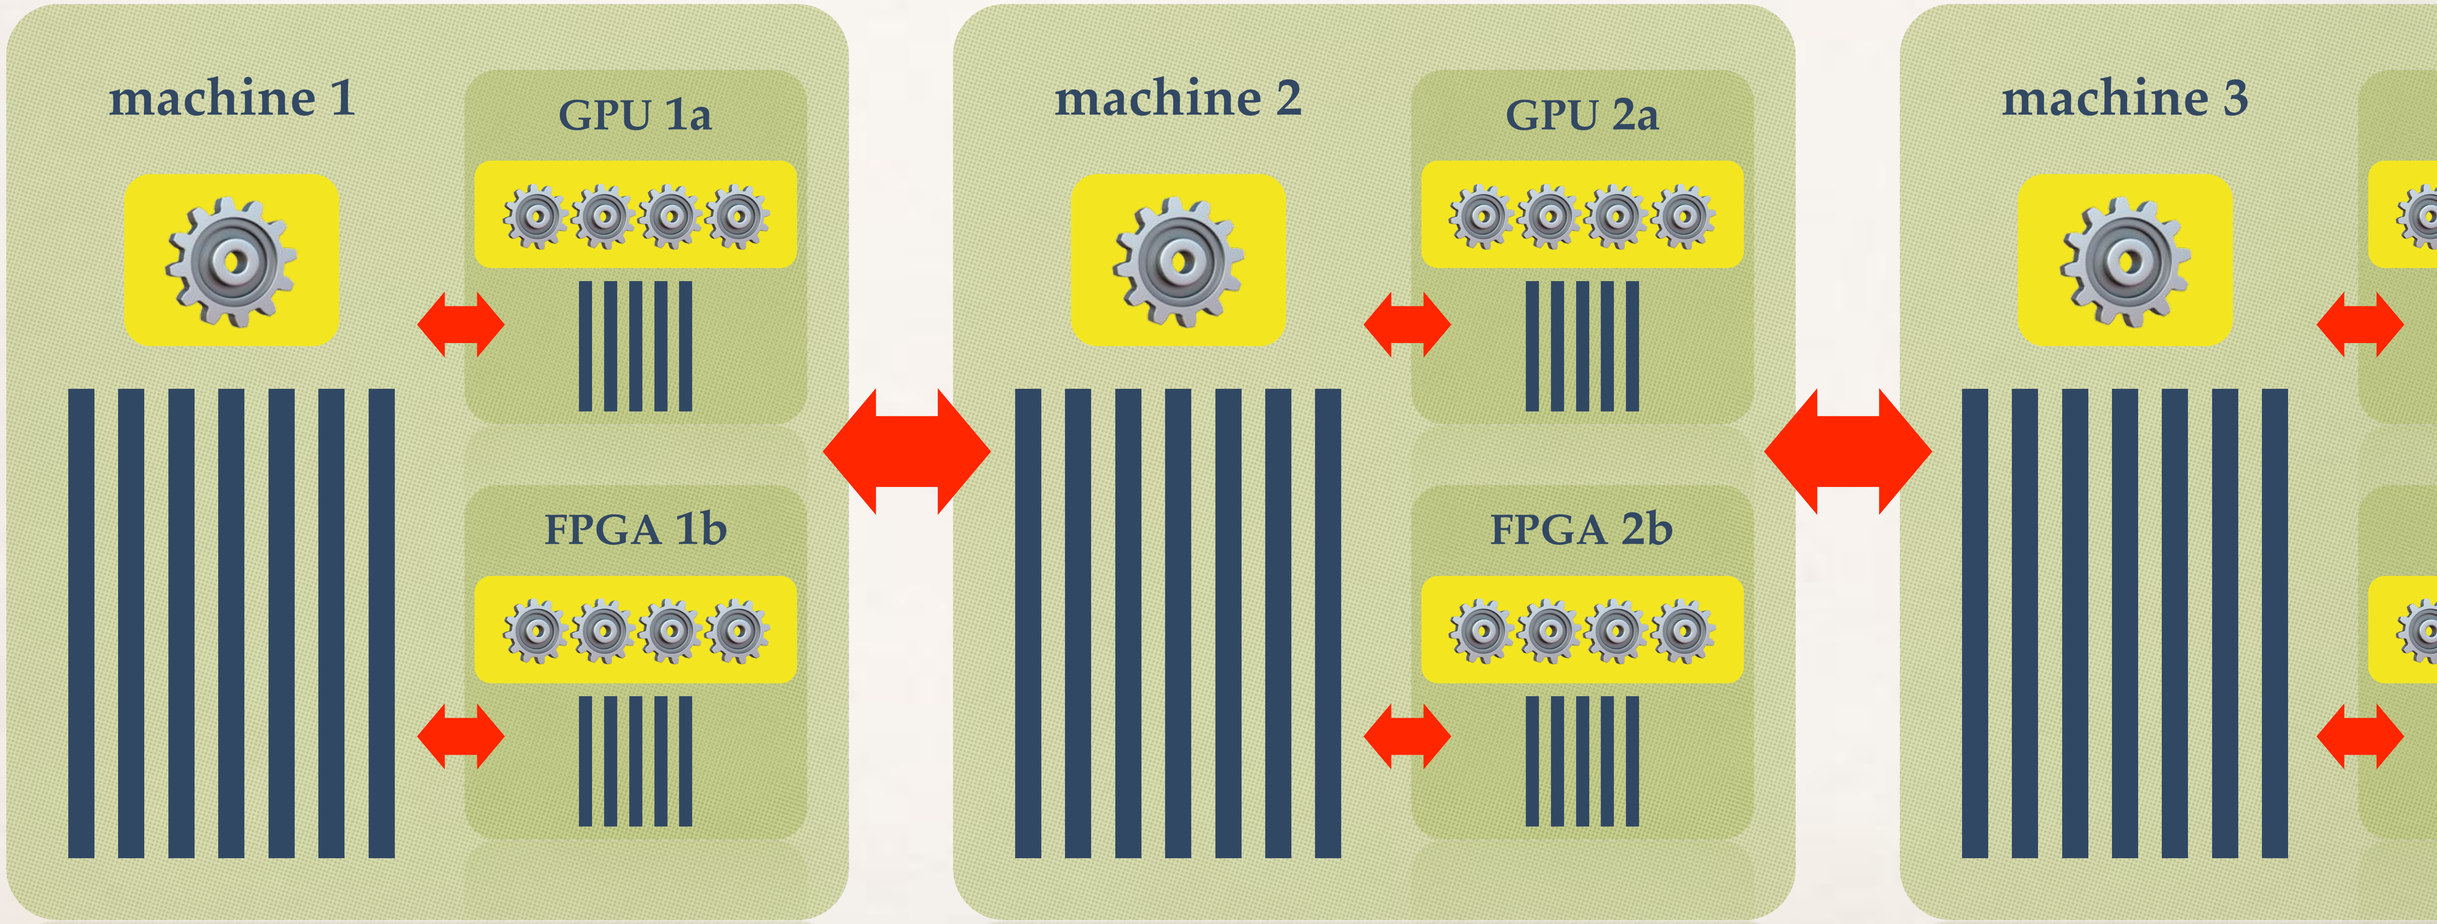

Challenge

The Cost of Communication

$\boldsymbol{v} \in \mathbb{R}^{100}$

- Reading $\boldsymbol{v}$ from memory (RAM)

  100 ns

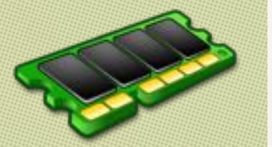

- Sending $\boldsymbol{v}$ to another machine

  500'000 ns

- Typical Map-Reduce iteration

  10'000'000'000 ns

Challenge

The Cost of Communication

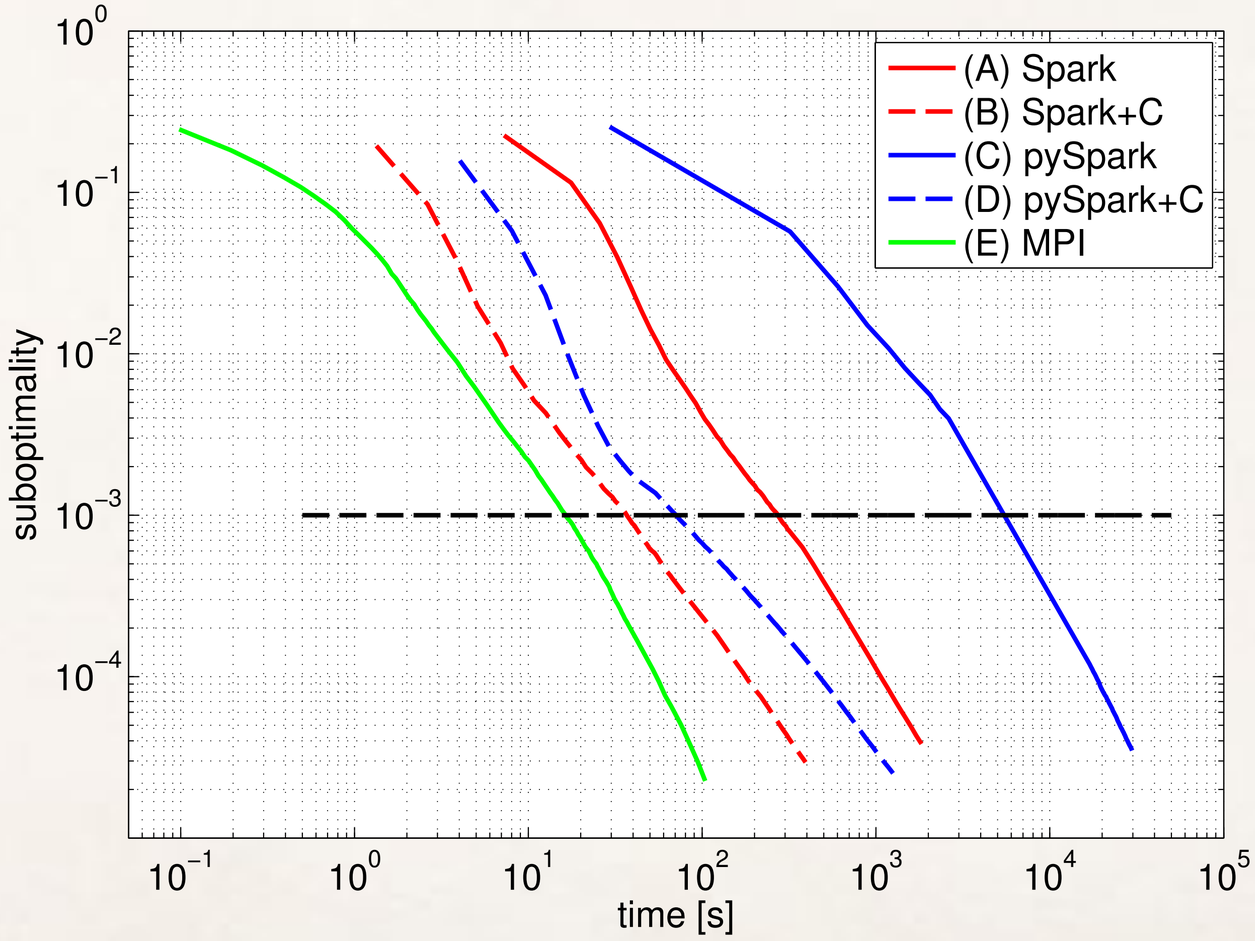

High-Performance Distributed Machine Learning using Apache Spark
Dünner et al. 2016, arxiv.org/abs/1612.01437

Challenge

Usability

Good distributed and parallel code is hard

- no **reusability** of good single machine algorithms & code
- no portability: model-specific and system-specific code

1

Distributed

What if the data does not fit onto one device anymore?

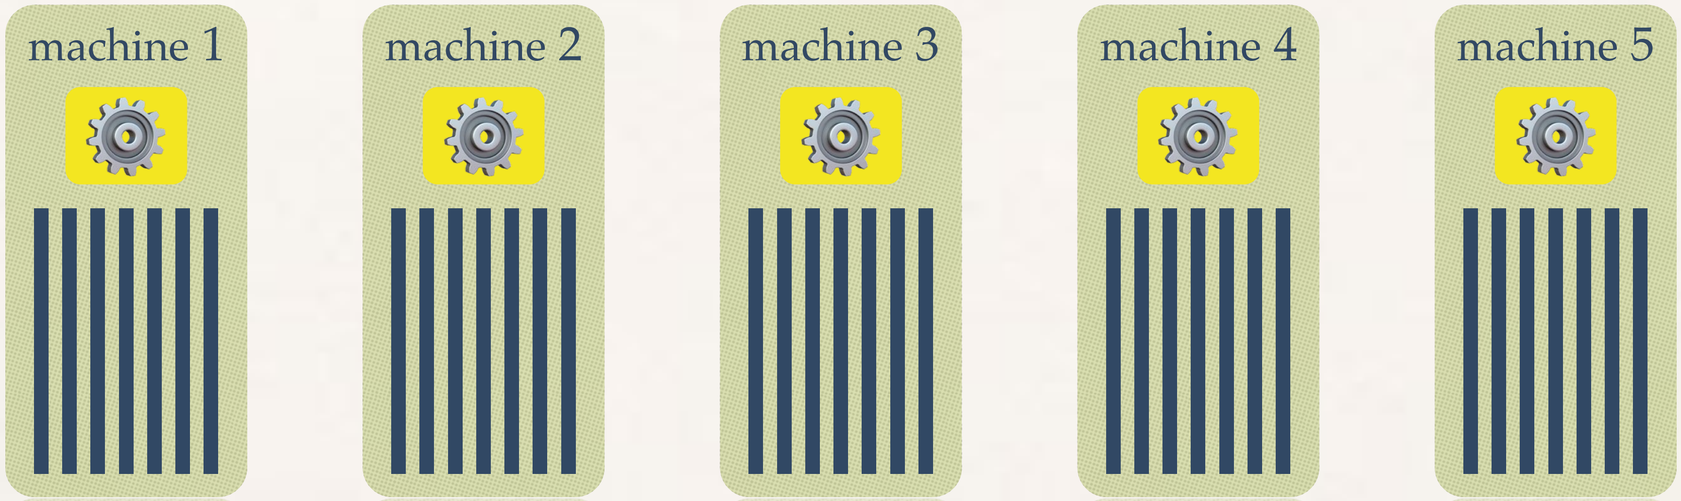

Problem class

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^n} \quad f(A\boldsymbol{\alpha}) \; + \; g(\boldsymbol{\alpha})$$

One-Shot Averaging Does Not Work

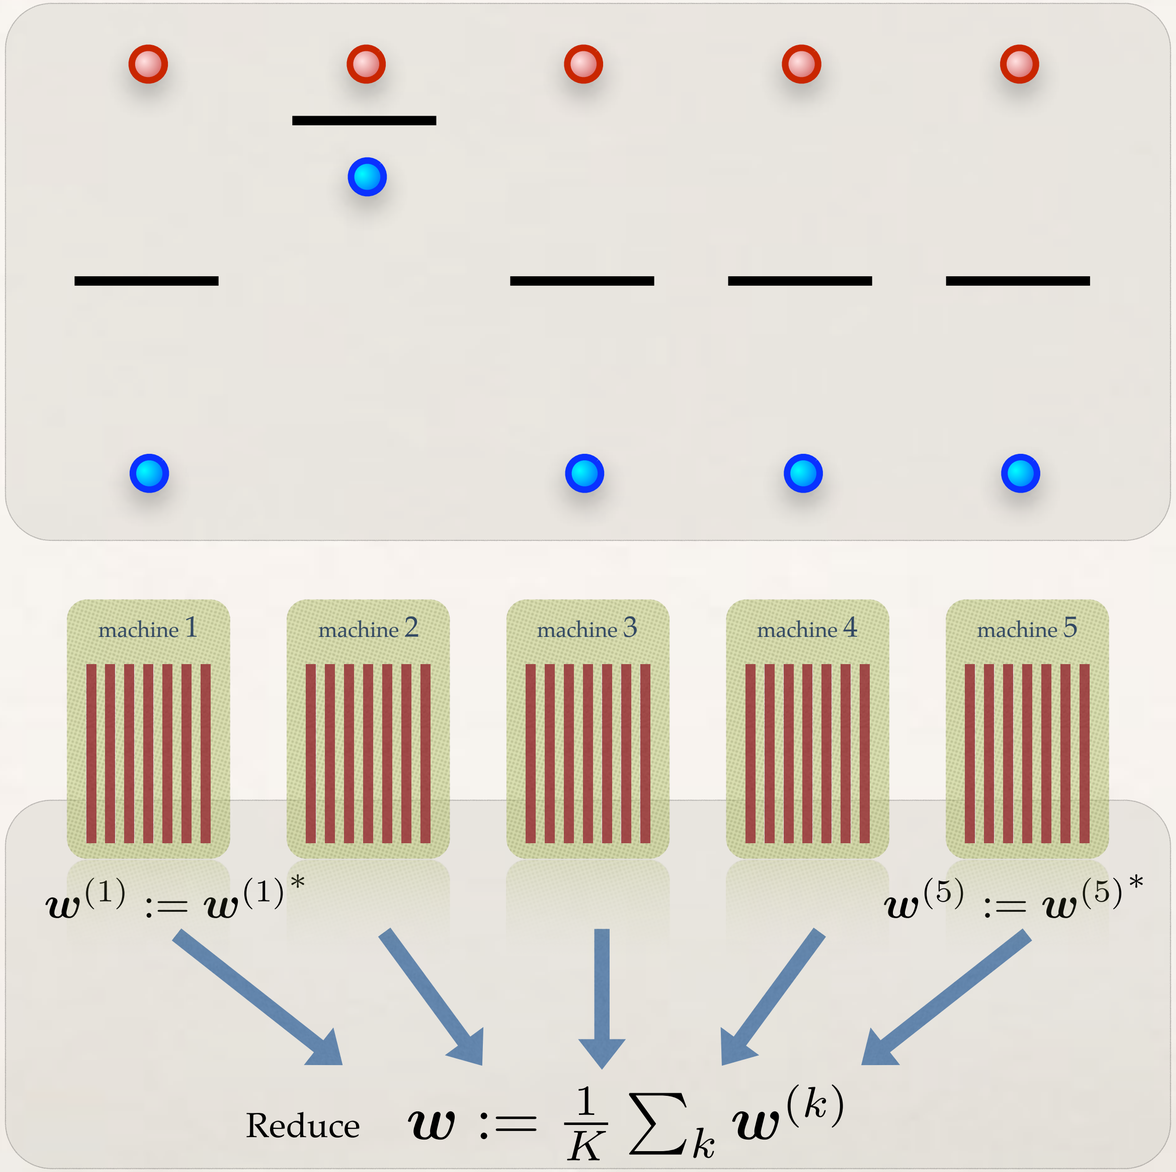

Optimization: Primal-Dual Context

$$A_{\text{loc}}\Delta\boldsymbol{\alpha}_{[k]} + \boldsymbol{w}$$

$$\min_{\boldsymbol{\alpha}\in\mathbb{R}^n} \quad \Big[\ \mathcal{O}_A(\boldsymbol{\alpha}) := f(A\boldsymbol{\alpha}) + g(\boldsymbol{\alpha})\ \Big]$$

primal Lasso
dual L2-reg SVM/Log-Regr
primal L1-reg SVM/Log-Reg

correspondence

$$\boldsymbol{w} := \nabla f(A\boldsymbol{\alpha})$$

$$\min_{\mathbf{w}\in\mathbb{R}^d} \quad \Big[\ \mathcal{O}_B(\mathbf{w}) := g^*(-A^\top\mathbf{w}) + f^*(\mathbf{w})\ \Big]$$

CoCoA - Communication Efficient Distributed Optimization

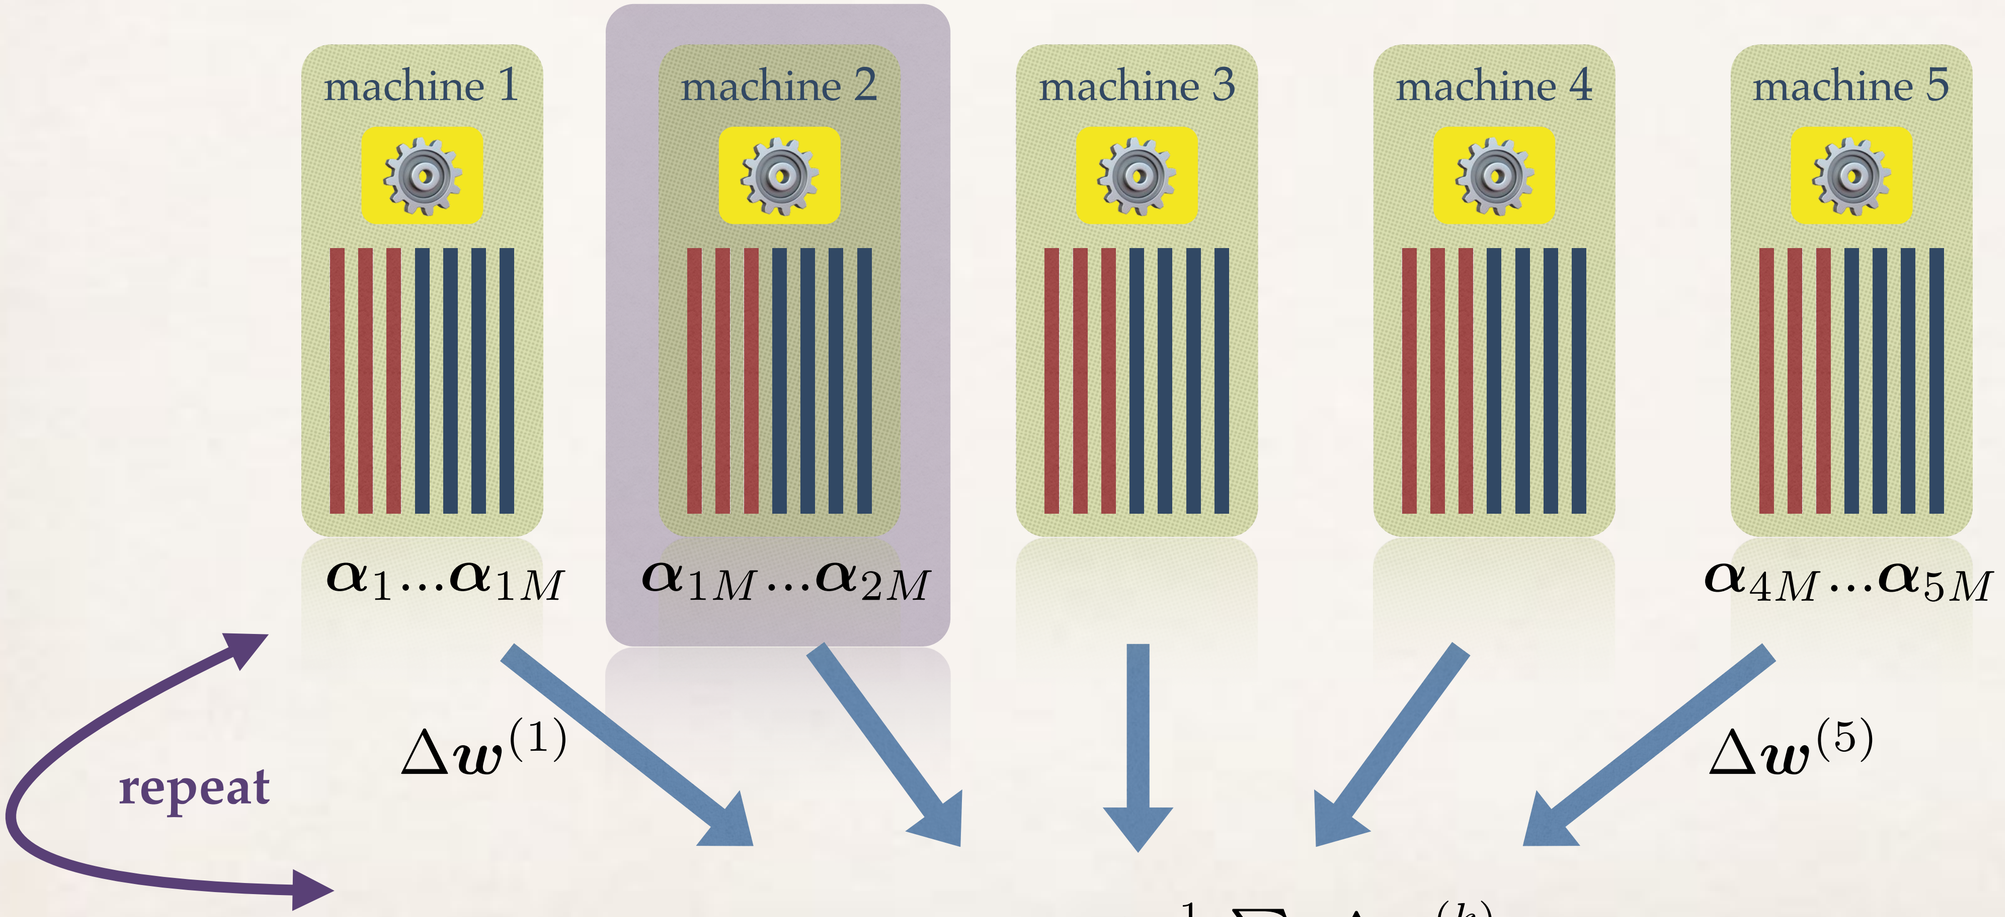

$$w := w + \frac{1}{K} \sum_k \Delta w^{(k)}$$

Distributed Experiments

Dataset	Training	Features	Sparsity
url	2,396,130	3,231,961	3.5e-3%
epsilon	400,000	2,000	100%
kddb	19,264,097	29,890,095	9.8e-5%
webspam	350,000	16,609,143	0.02%

L1-Regularized Linear Regression

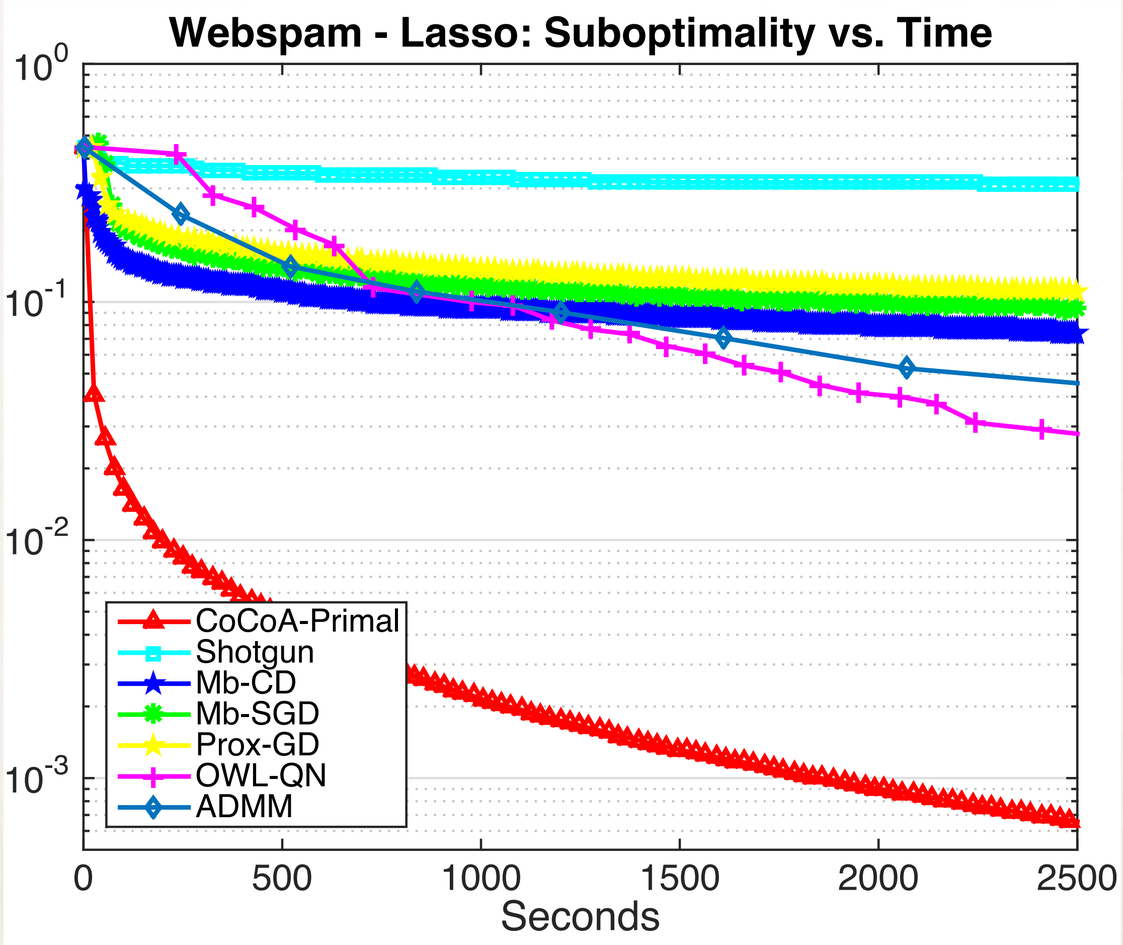

NIPS 2014, ICML 2015,
arxiv.org/abs/1611.02189

- part of TensorFlow core (L2)
- *custom code (L1), TF,spark,C*

Summary

- **adaptivity** to the communication cost
- **re-usability** of good existing solvers
- **accuracy** certificates

Next Steps

- **second-order** and **trust-region** version (local Hessian)
- **adaptivity** to the degree of separability *(coming soon)*
- generalization to **deep learning, SGD**
- **benchmarking** & code

2

Leveraging Heterogenous Systems

Compute & Memory Hierarchy: Which data to put in which device?

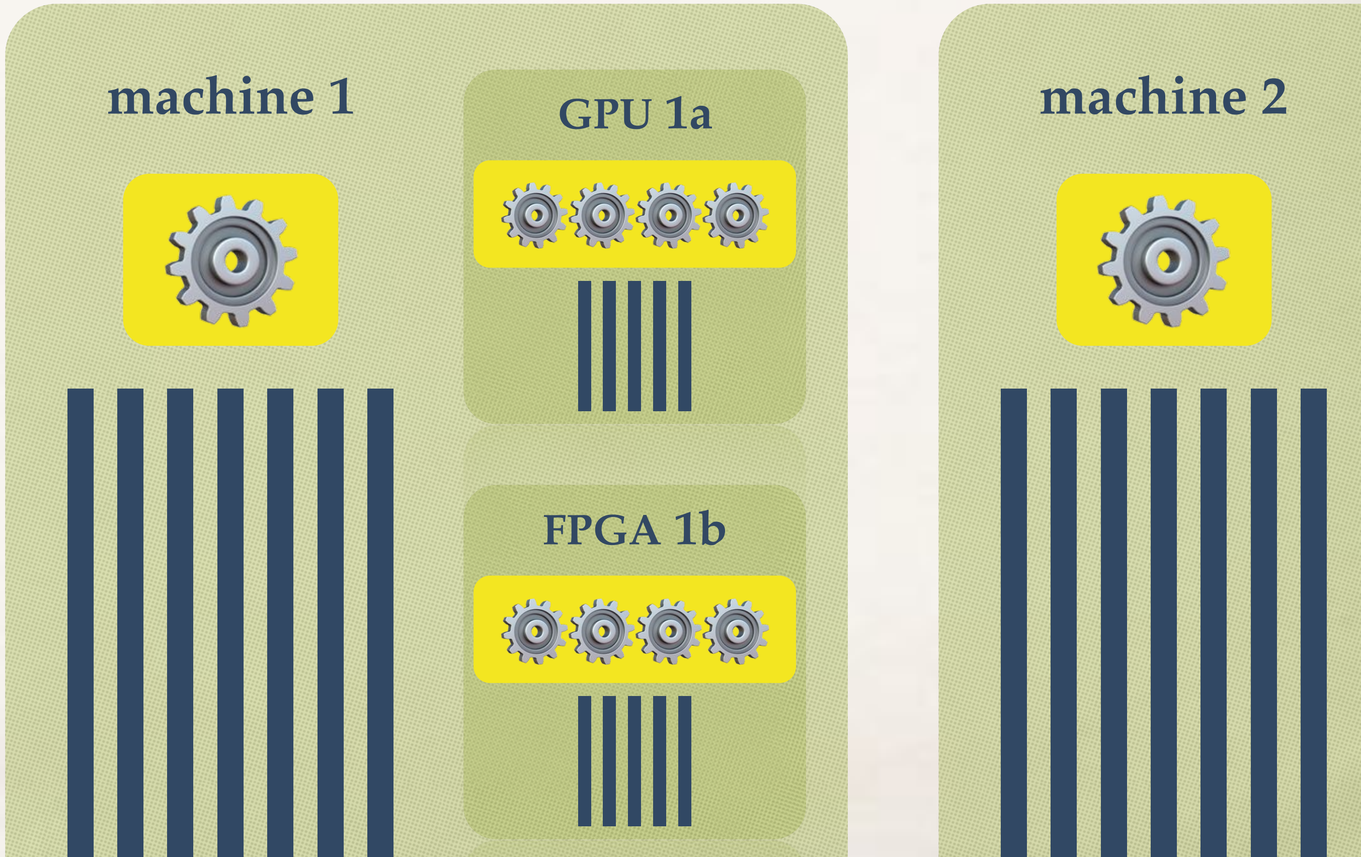

Leveraging Heterogenous Systems

duality gap as selection criterion

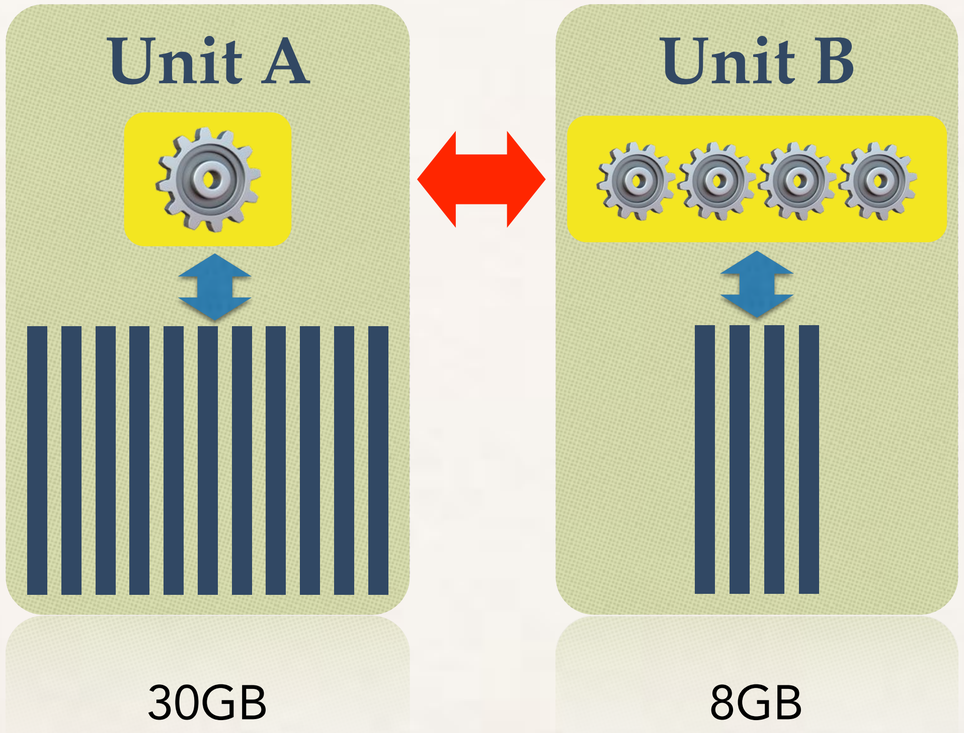

adaptive importance sampling

AISTATS 2017, 2018
NIPS 2017a,b

Experiments

RAM ⟷ GPU, 30GB dataset

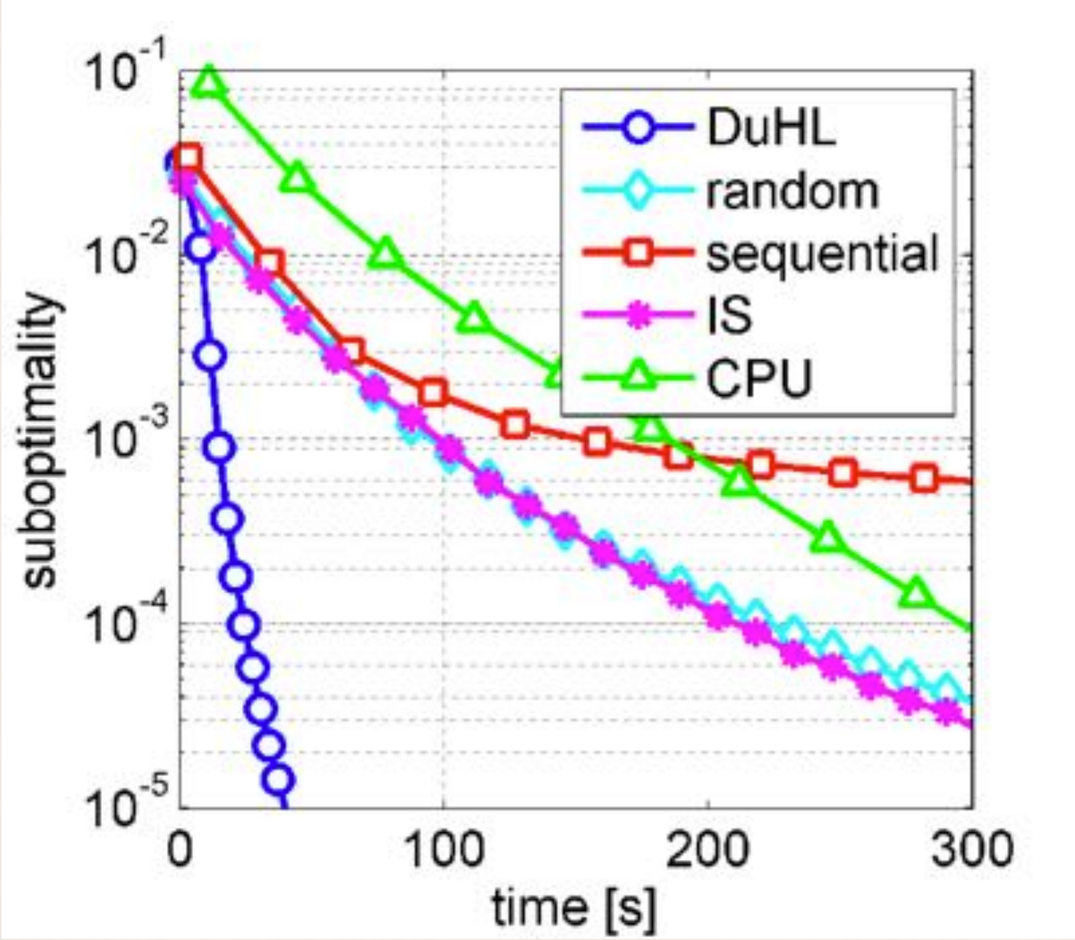

Lasso

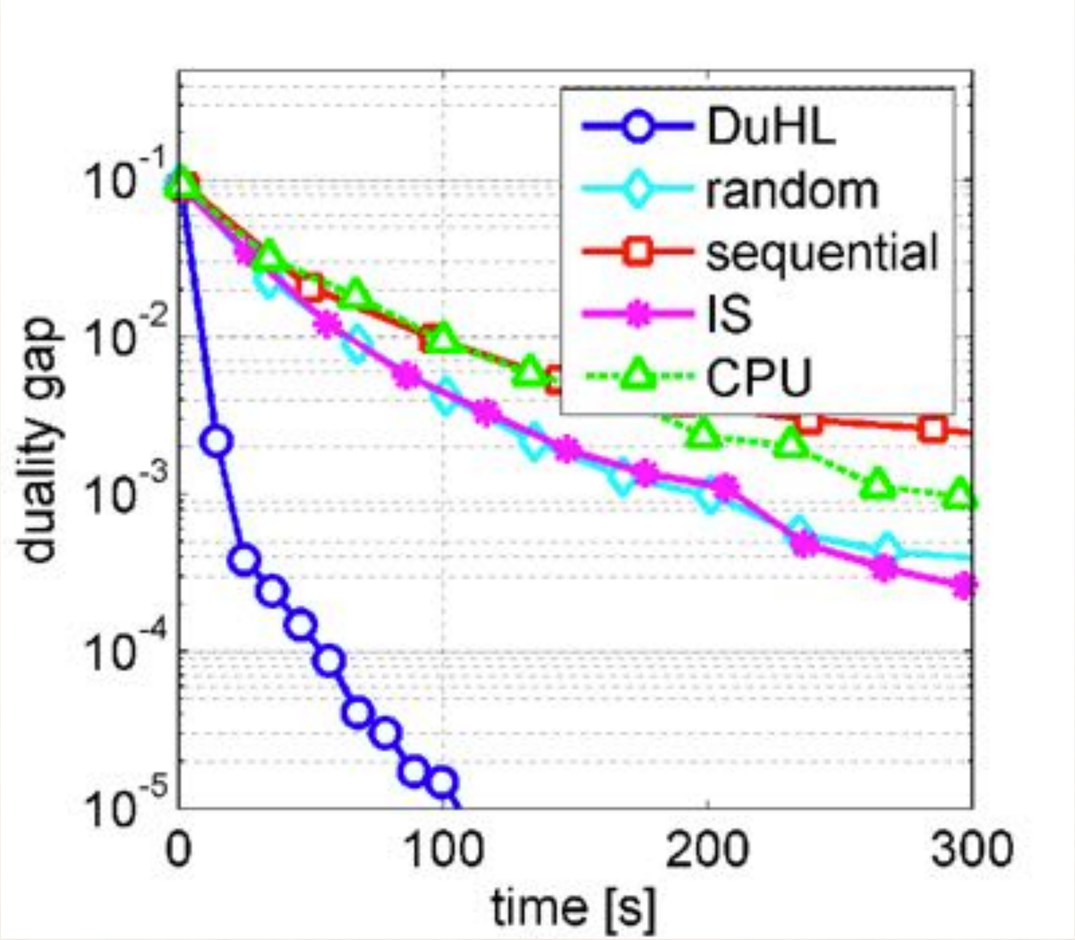

SVM

Open Research

- **limited precision operations** for efficiency of communication and computation
- **asynchronous** and **fault tolerant** algorithms
- **multi-level approach** on heterogenous systems
- more **re-usable** algorithmic building blocks
  - for more systems and problems

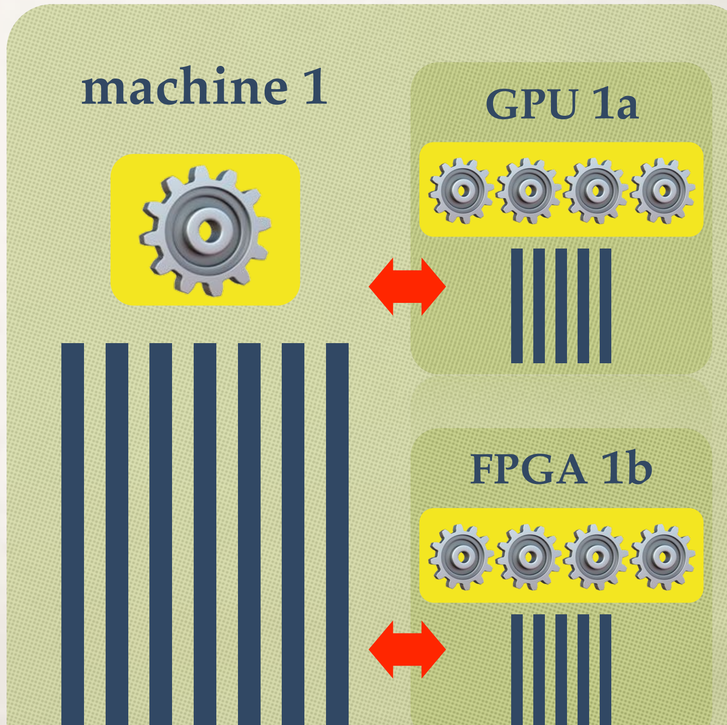

Thanks!

mlo.epfl.ch

Celestine Dünner, Virginia Smith, Simone Forte, Thomas Parnell, Chenxin Ma, Martin Takac, Dmytro Perekrestenko, Volkan Cevher, Michael I. Jordan, Thomas Hofmann, Sebastian Stich, Anant Raj