The Berlin Big Data Center

Jonas Traub

Technische Universität Berlin / DFKI IAM

www.dima.tu-berlin.de | bbdc.berlin | jonas.traub@tu-berlin.de

with materials from Tilmann Rabl and Volker Markl

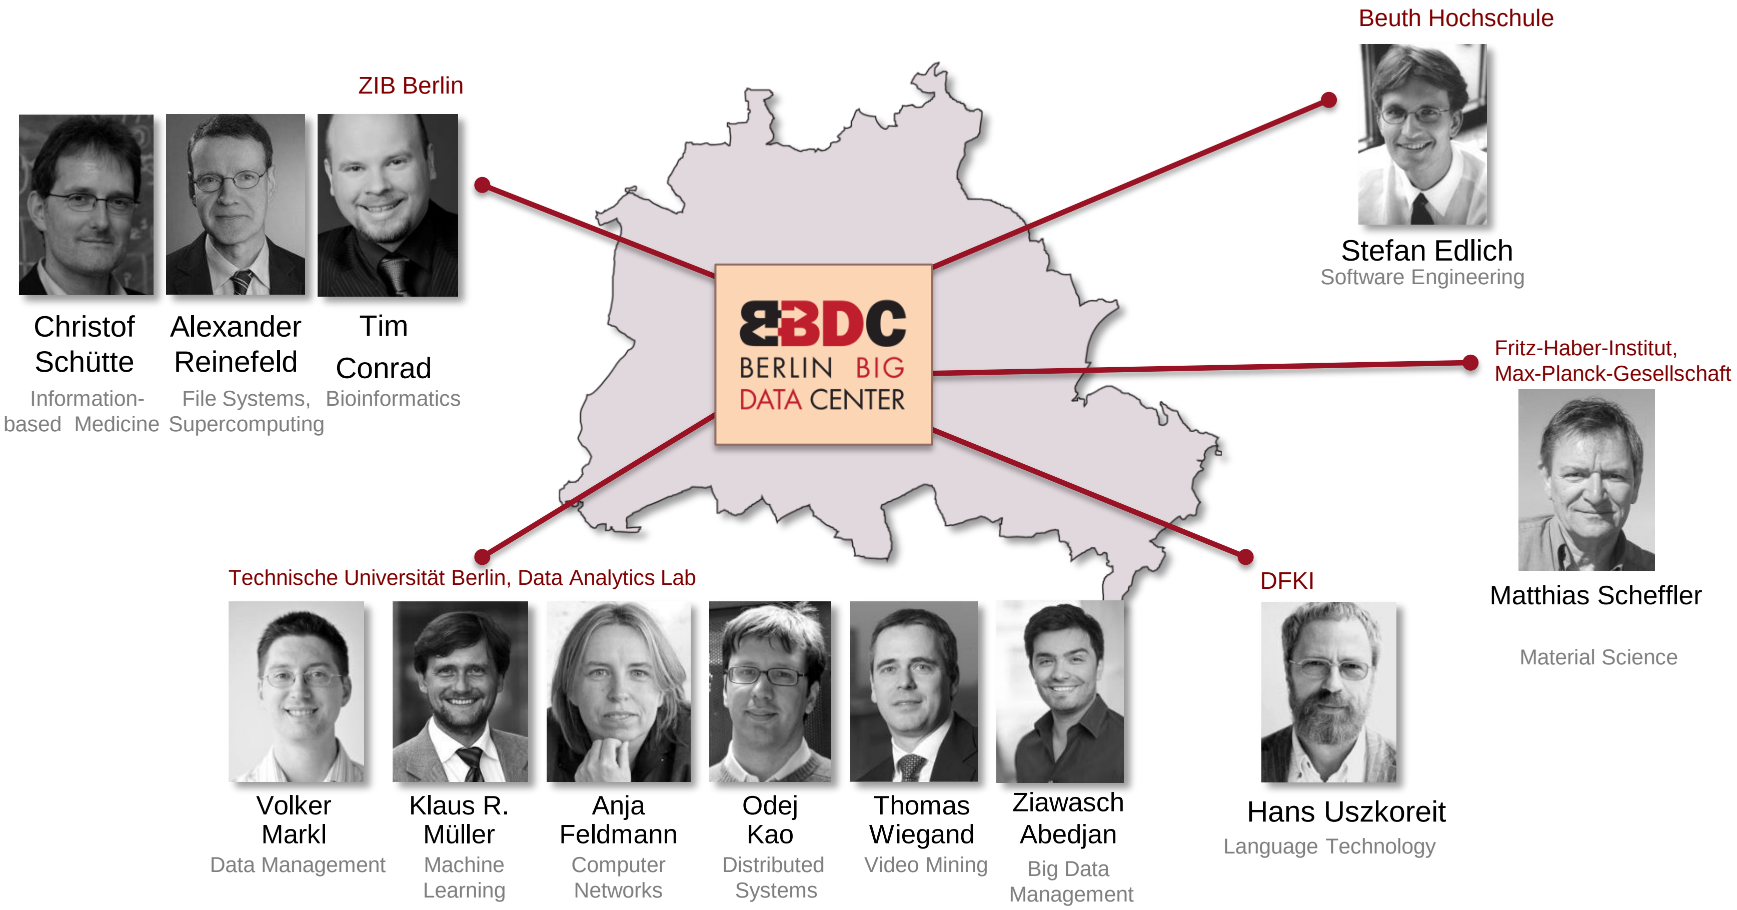

BBDC
BERLIN BIG DATA CENTER

ZIB Berlin

Christof Schütte
Information-based Medicine

Alexander Reinefeld
File Systems, Supercomputing

Tim Conrad
Bioinformatics

Beuth Hochschule

Stefan Edlich
Software Engineering

Fritz-Haber-Institut,
Max-Planck-Gesellschaft

Matthias Scheffler
Material Science

Technische Universität Berlin, Data Analytics Lab

Volker Markl
Data Management

Klaus R. Müller
Machine Learning

Anja Feldmann
Computer Networks

Odej Kao
Distributed Systems

Thomas Wiegand
Video Mining

Ziawasch Abedjan
Big Data Management

DFKI

Hans Uszkoreit
Language Technology

BBDC Goals

- **Pooling expertise** in scalable data management, data analytics, and big data applications.
- **Conducting fundamental research** to develop novel and automatically scalable technologies capable of performing “deep analysis” of big data.
- **Developing an integrated, declarative, highly scalable, open-source system** that enables the specification, automatic optimization, parallelization, hardware adaptation, fault-tolerance, and efficient execution of advanced data analysis problems using varying methods that leverage our work on Apache Flink.
- **Transferring technology and know-how** to support innovation in companies and startups.
- **Educating data scientists** with respect to the five big data dimensions via leading educational programs.
- **Empowering people to leverage** “Smart Data.”
- **Enabling the general public to conduct sound data-driven decision-making.**

Big Data – System View

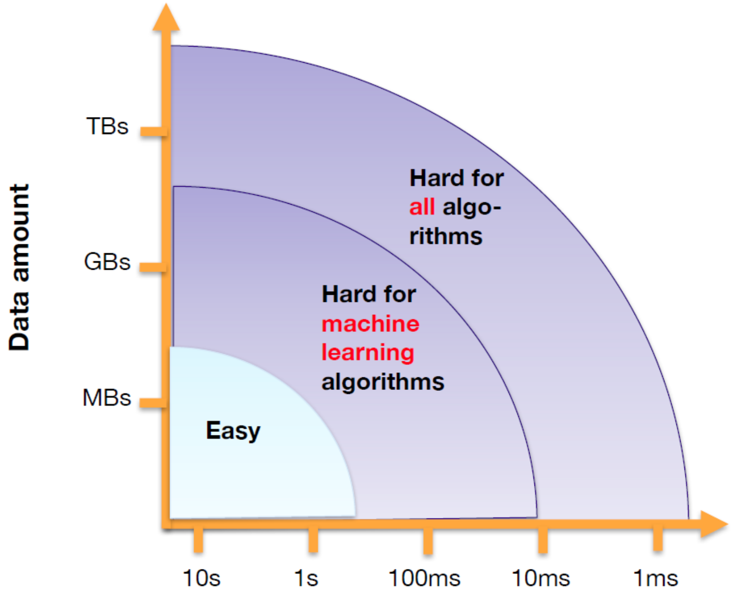

Tension between *performance* and *algorithmic expressiveness*

“Data Scientist” – “Jack of All Trades!”

Domain Expertise (e.g., Industry 4.0, Medicine, Physics, Engineering, Energy, Logistics)

Mathematical Programming

Linear Algebra

Stochastic Gradient Descent

Error Estimation

Active Sampling

Regression

Monte Carlo

Statistics

Sketches

Hashing

Convergence

Decoupling

Iterative Algorithms

Curse of Dimensionality

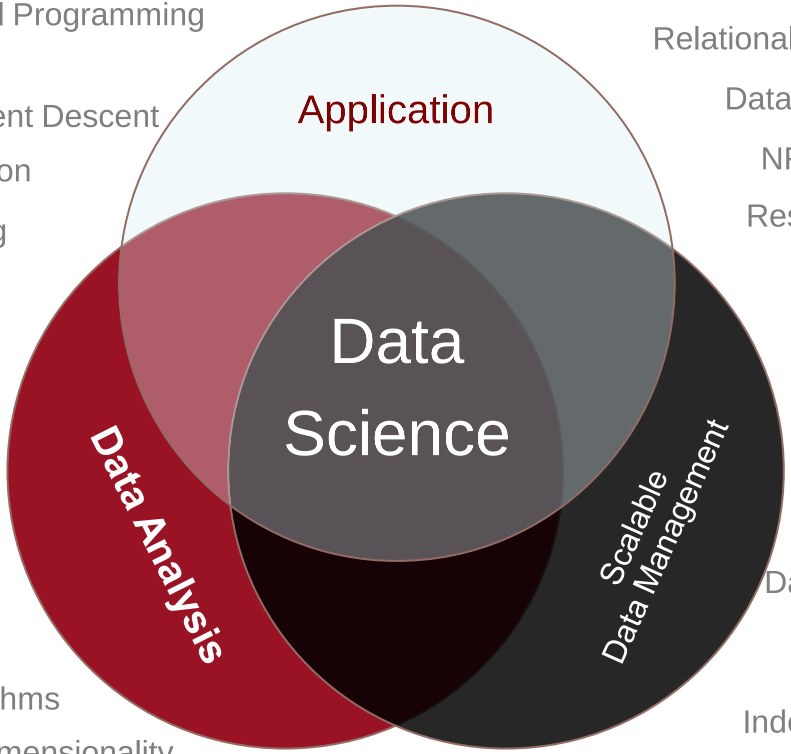

Relational Algebra / SQL

Data Warehouse/OLAP

NF^2 /XQuery

Resource Management

Hardware Adaptation

Fault Tolerance

Memory Management

Parallelization

Scalability

Memory Hierarchy

Data Analysis Language Compiler

Query Optimization

Indexing

Data Flow

Control Flow

Real-Time

Big Data Analytics Requires Systems Programming

Data Analysis
Statistics
Algebra
Optimization
Machine Learning
NLP
Signal Processing
Image Analysis
Audio-,Video Analysis
Information Integration
Information Extraction
Data Value Chain
Data Analysis Process
Predictive Analytics

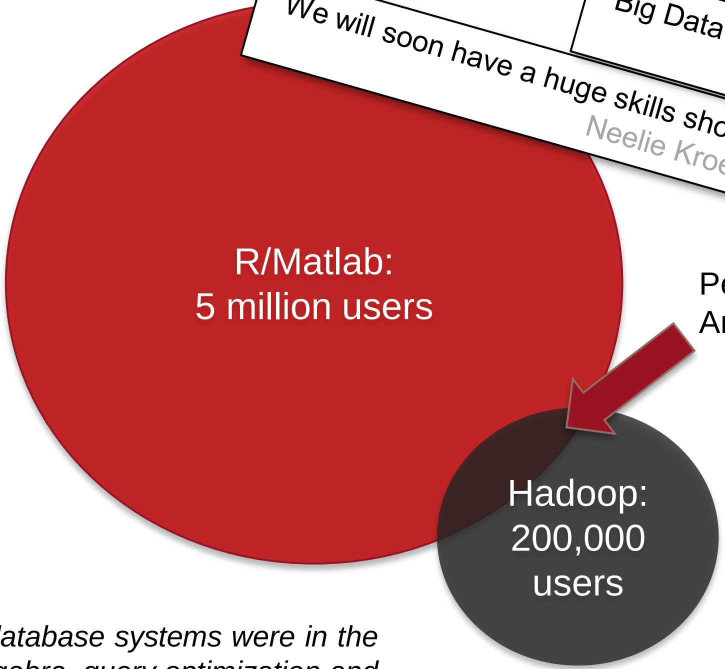

“Big Data‘s Big Problem: Little Talent“
Wall Street Journal

“We will soon have a huge skills shortage for data-related jobs.“
Neelie Kroes (ICT 2013, Nov. 7, Vilnius)

People with Big Data Analytics Skills

Indexing
Parallelization
Communication
Memory Management
Query Optimization
Efficient Algorithms
Resource Management
Fault Tolerance
Numerical Stability

Big Data is now where database systems were in the 70s (prior to relational algebra, query optimization and a SQL-standard)!

Declarative languages to the rescue!

“Data Scientist” – “Jack of All Trades!”

Domain Expertise (e.g., Industry 4.0, Medicine, Physics, Engineering, Energy, Logistics)

Mathematical Programming

Linear Algebra

Stochastic Gradient Descent

Error Estimation

Active Sampling

Regression

Monte Carlo

Statistics

Sketches

Hashing

Convergence

Decoupling

Iterative Algorithms

Curse of Dimensionality

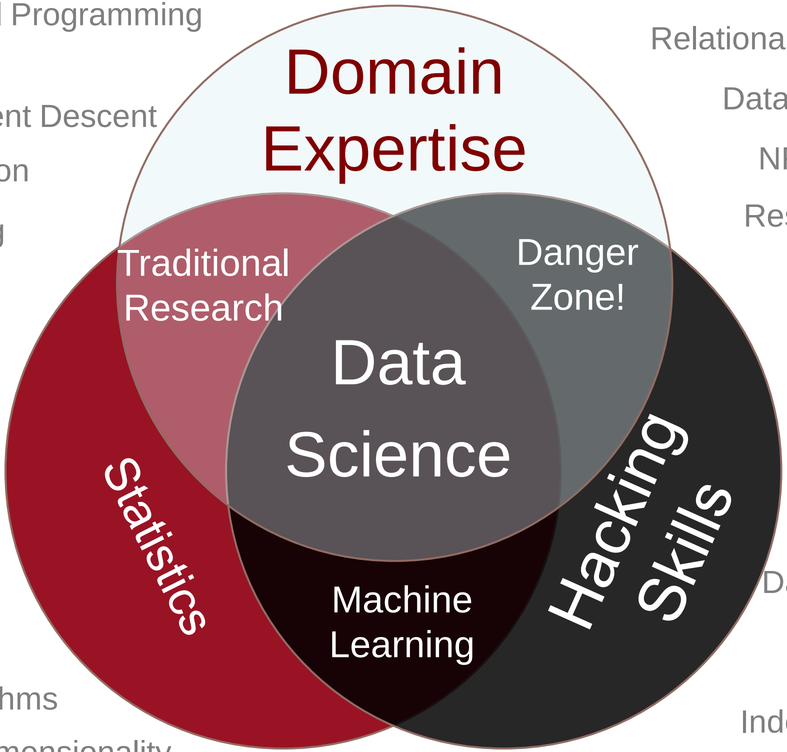

Relational Algebra / SQL

Data Warehouse/OLAP

NF^2 /XQuery

Resource Management

Hardware Adaptation

Fault Tolerance

Memory Management

Parallelization

Scalability

Memory Hierarchy

Data Analysis Language
Compiler

Query Optimization

Indexing

Data Flow

Control Flow

Real-Time

New Technology to the Rescue!

Danger Zone!

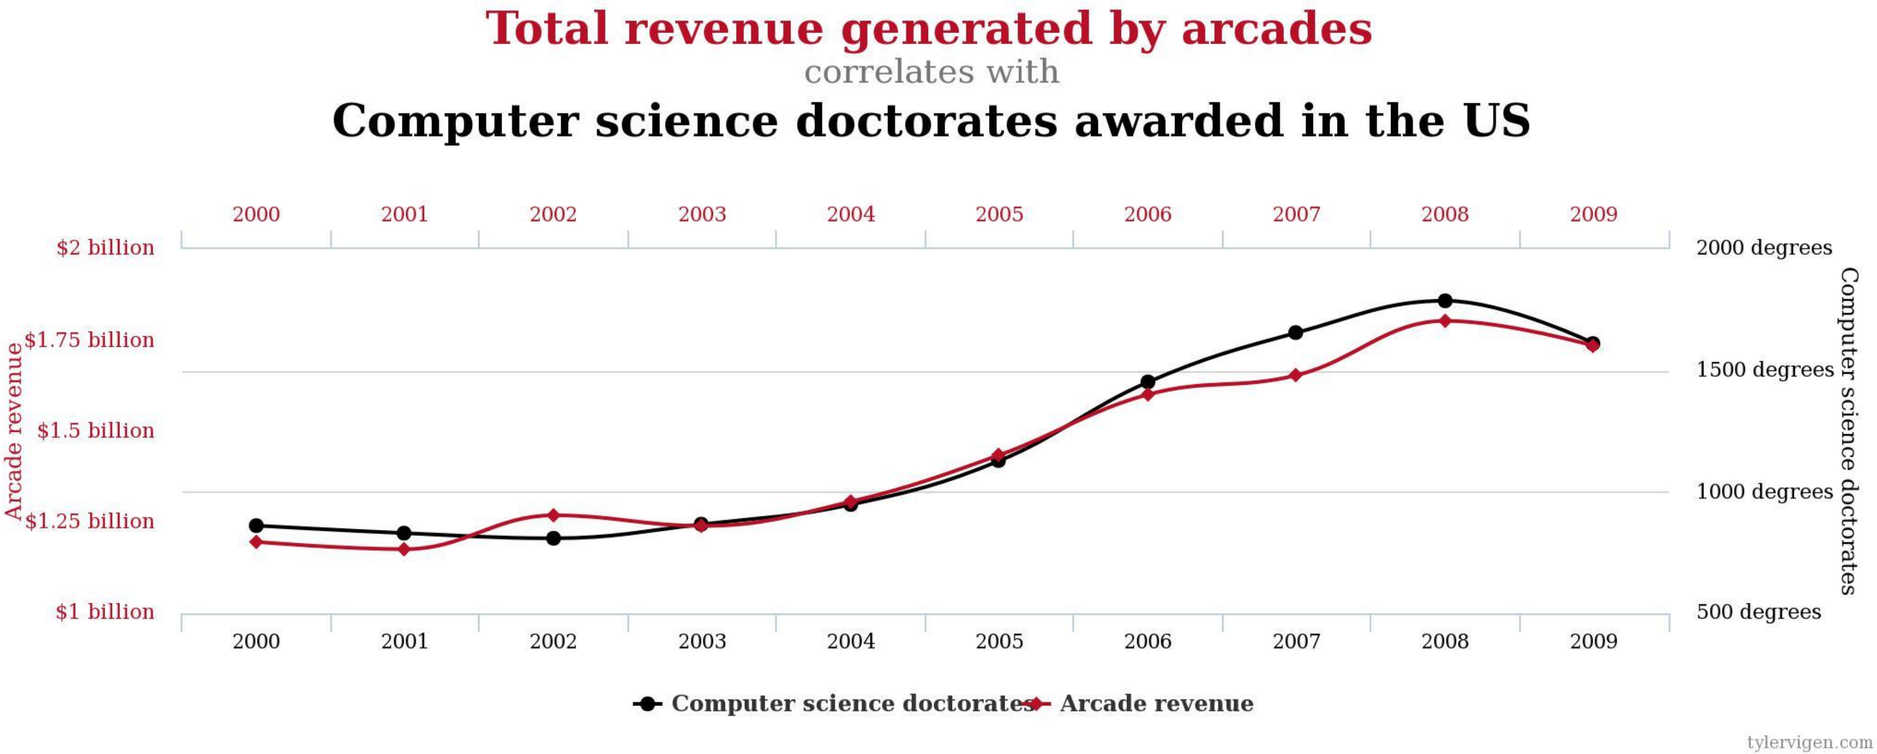
Total revenue generated by arcades
correlates with
Computer science doctorates awarded in the US
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
$2 billion
$1.75 billion
$1.5 billion
$1.25 billion
$1 billion
Arcade revenue
2000 degrees
1500 degrees
1000 degrees
500 degrees
Computer science doctorates
Computer science doctorates
Arcade revenue
tylervigen.com

Danger Zone! Contd.

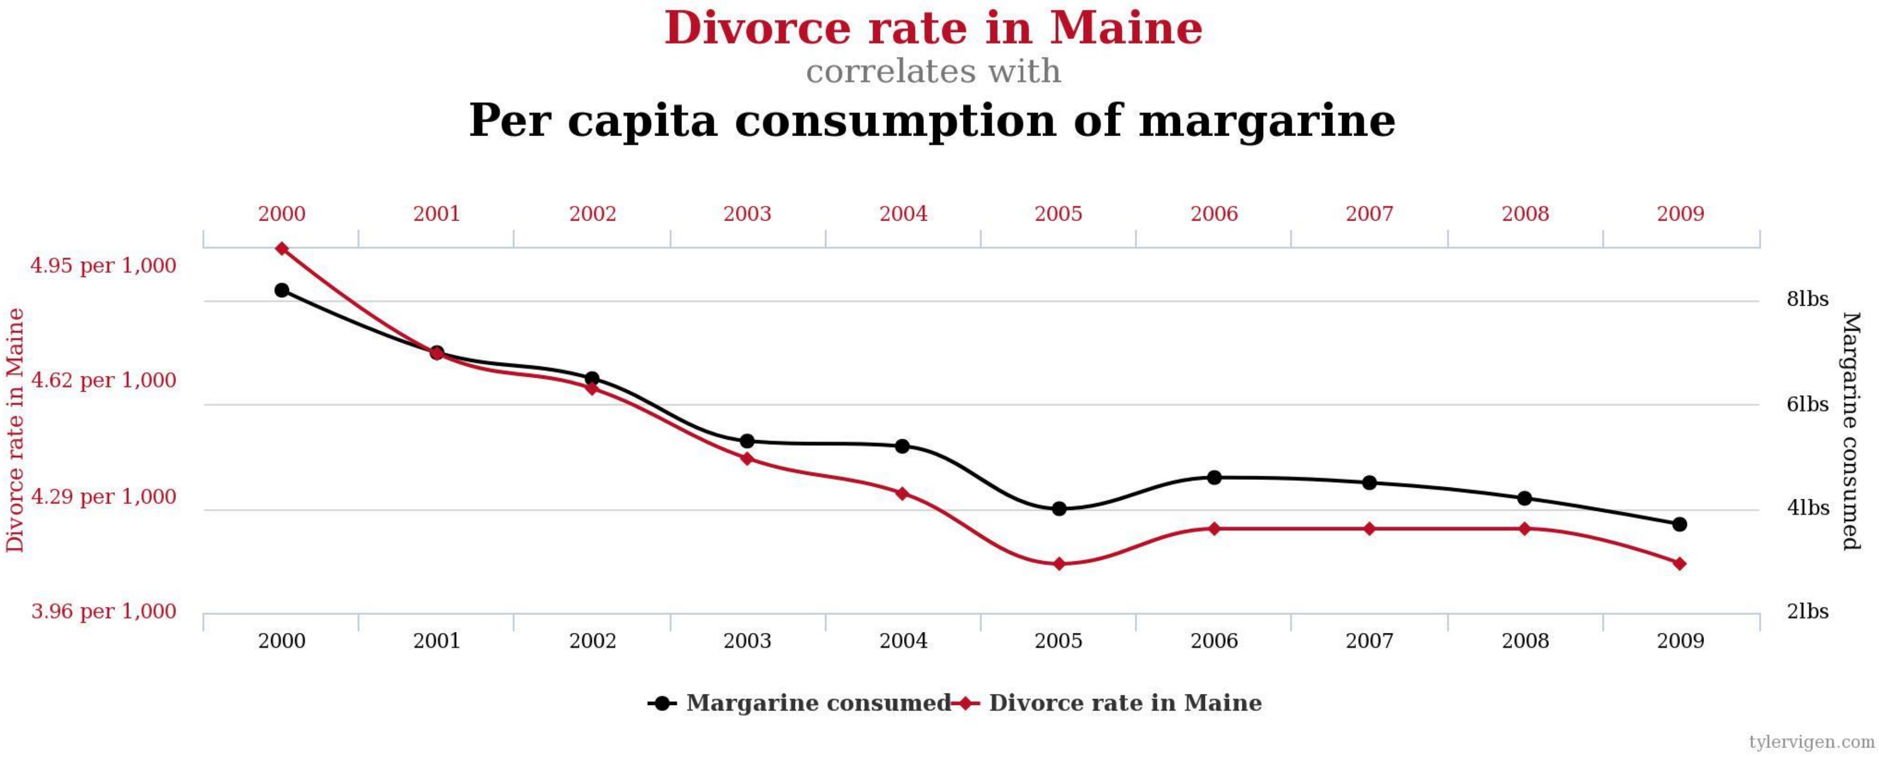

Apache Flink–
A Success Story
Born in Berlin

http://flink.apache.org

Timeline

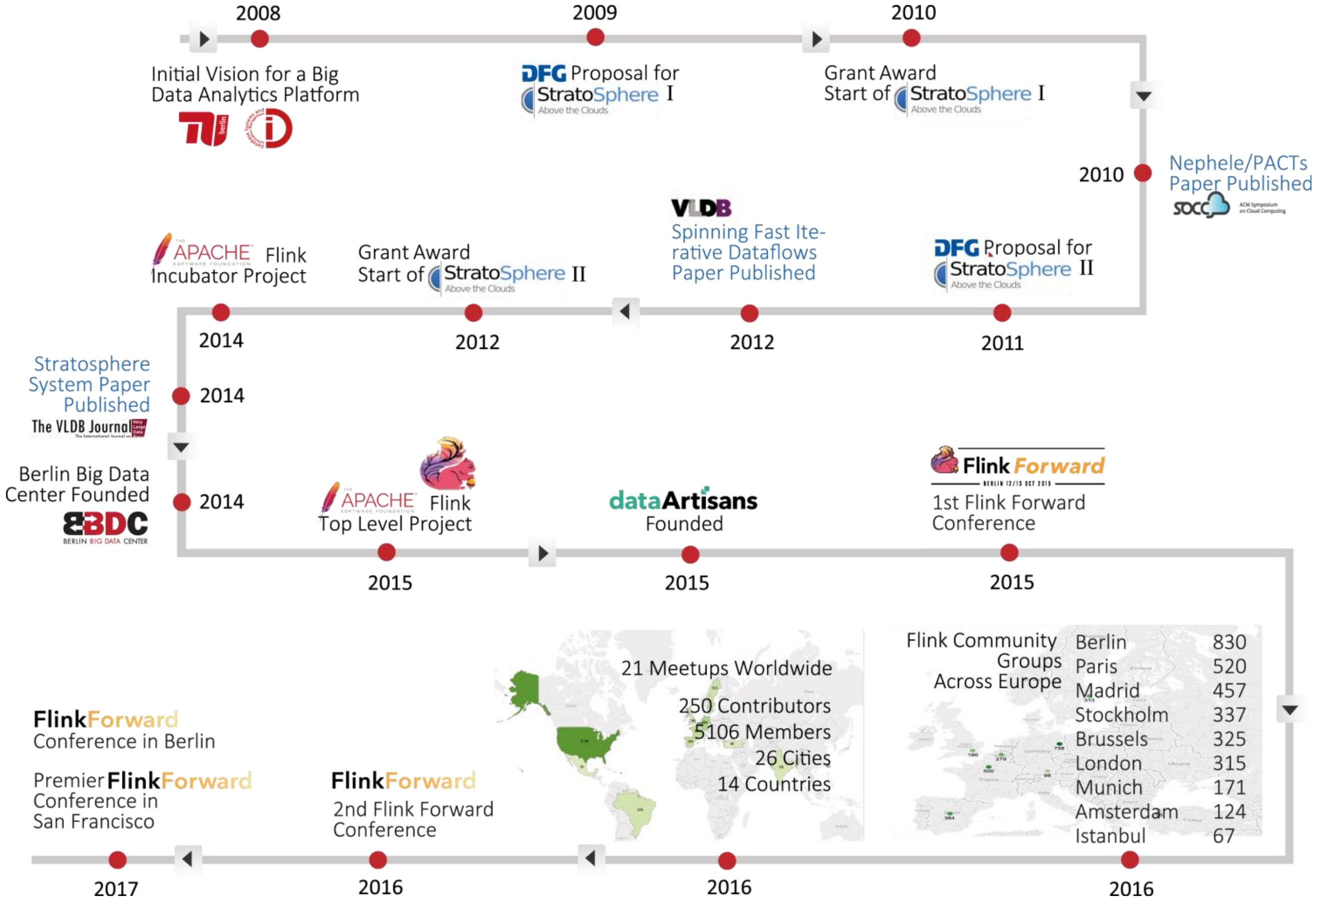

2008
Initial Vision for a Big Data Analytics Platform

2009
DFG Proposal for StratoSphere I
Above the Clouds

2010
Grant Award
Start of StratoSphere I
Above the Clouds

2010
Nephele/PACTs Paper Published
SoCC ACM Symposium on Cloud Computing

2011
DFG Proposal for StratoSphere II
Above the Clouds

2012
VLDB
Spinning Fast Iterative Dataflows Paper Published

2012
Grant Award
Start of StratoSphere II
Above the Clouds

2014
APACHE Flink Incubator Project

2014
Stratosphere System Paper Published
The VLDB Journal

2014
Berlin Big Data Center Founded
BBDC BERLIN BIG DATA CENTER

2015
APACHE Flink Top Level Project

2015
dataArtisans Founded

2015
Flink Forward BERLIN 12/13 OCT 2015
1st Flink Forward Conference

2016
Flink Community Groups Across Europe

Berlin	830
Paris	520
Madrid	457
Stockholm	337
Brussels	325
London	315
Munich	171
Amsterdam	124
Istanbul	67

2016
21 Meetups Worldwide
250 Contributors
5106 Members
26 Cities
14 Countries

2016
FlinkForward
2nd Flink Forward Conference

2017
FlinkForward
Conference in Berlin

Premier FlinkForward
Conference in San Francisco

What can I do with it?

Batch processing

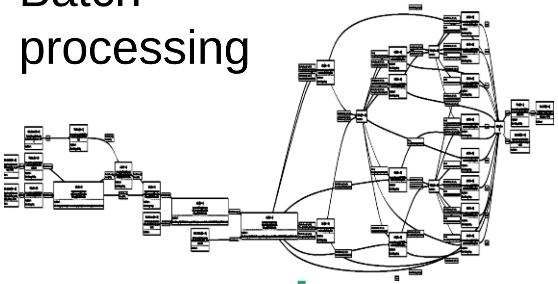

Machine Learning at scale

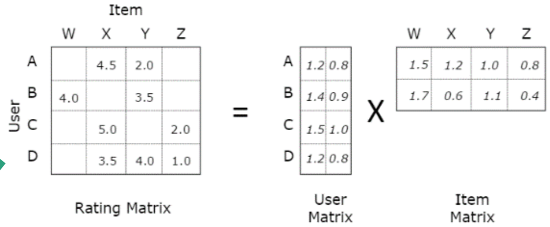

	Item W	X	Y	Z
User A		4.5	2.0	
B	4.0		3.5	
C		5.0		2.0
D		3.5	4.0	1.0

Rating Matrix

=

A	1.2	0.8
B	1.4	0.9
C	1.5	1.0
D	1.2	0.8

User Matrix

X

W	X	Y	Z
1.5	1.2	1.0	0.8
1.7	0.6	1.1	0.4

Item Matrix

Stream processing

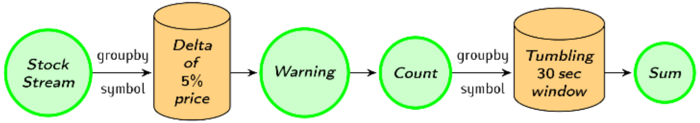

Graph Analysis

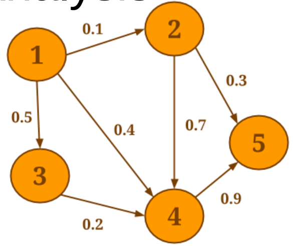

Flink

A big data processing system that can **natively** support all these workloads.

What is Apache Flink?

Apache Flink® is an open-source stream processing framework for distributed, high-performing, always-available, and accurate data streaming applications.

Key Features

- Bounded and unbounded data
- Event time semantics
- Stateful and fault-tolerant
- Running on thousands of nodes with very good throughput and latency
- Exactly-once semantics for stateful computations
- Flexible windowing based on time, count, or sessions in addition to data-driven windows
- **DataSet** and **DataStream** programming abstractions are the foundation for user programs and higher layers

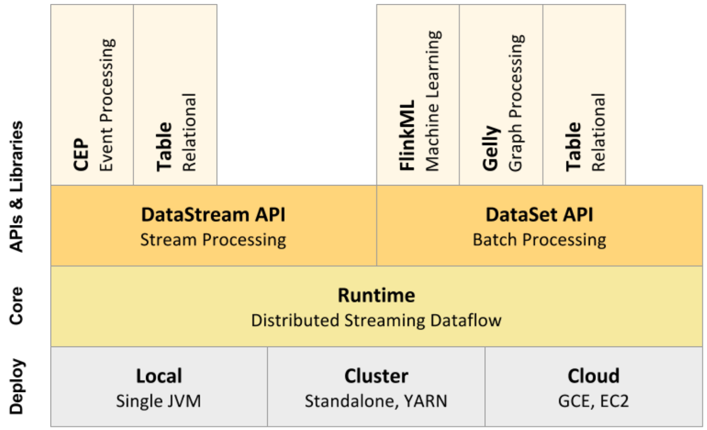

APIs & Libraries

CEP Event Processing

Table Relational

FlinkML Machine Learning

Gelly Graph Processing

Table Relational

DataStream API Stream Processing

DataSet API Batch Processing

Core

Runtime Distributed Streaming Dataflow

Deploy

Local Single JVM

Cluster Standalone, YARN

Cloud GCE, EC2

P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, K. Tzoumas:
Apache Flink™: Stream and Batch Processing in a Single Engine.
IEEE Data Eng. Bull. 38(4): 28-38 (2015)

Companies Using Flink

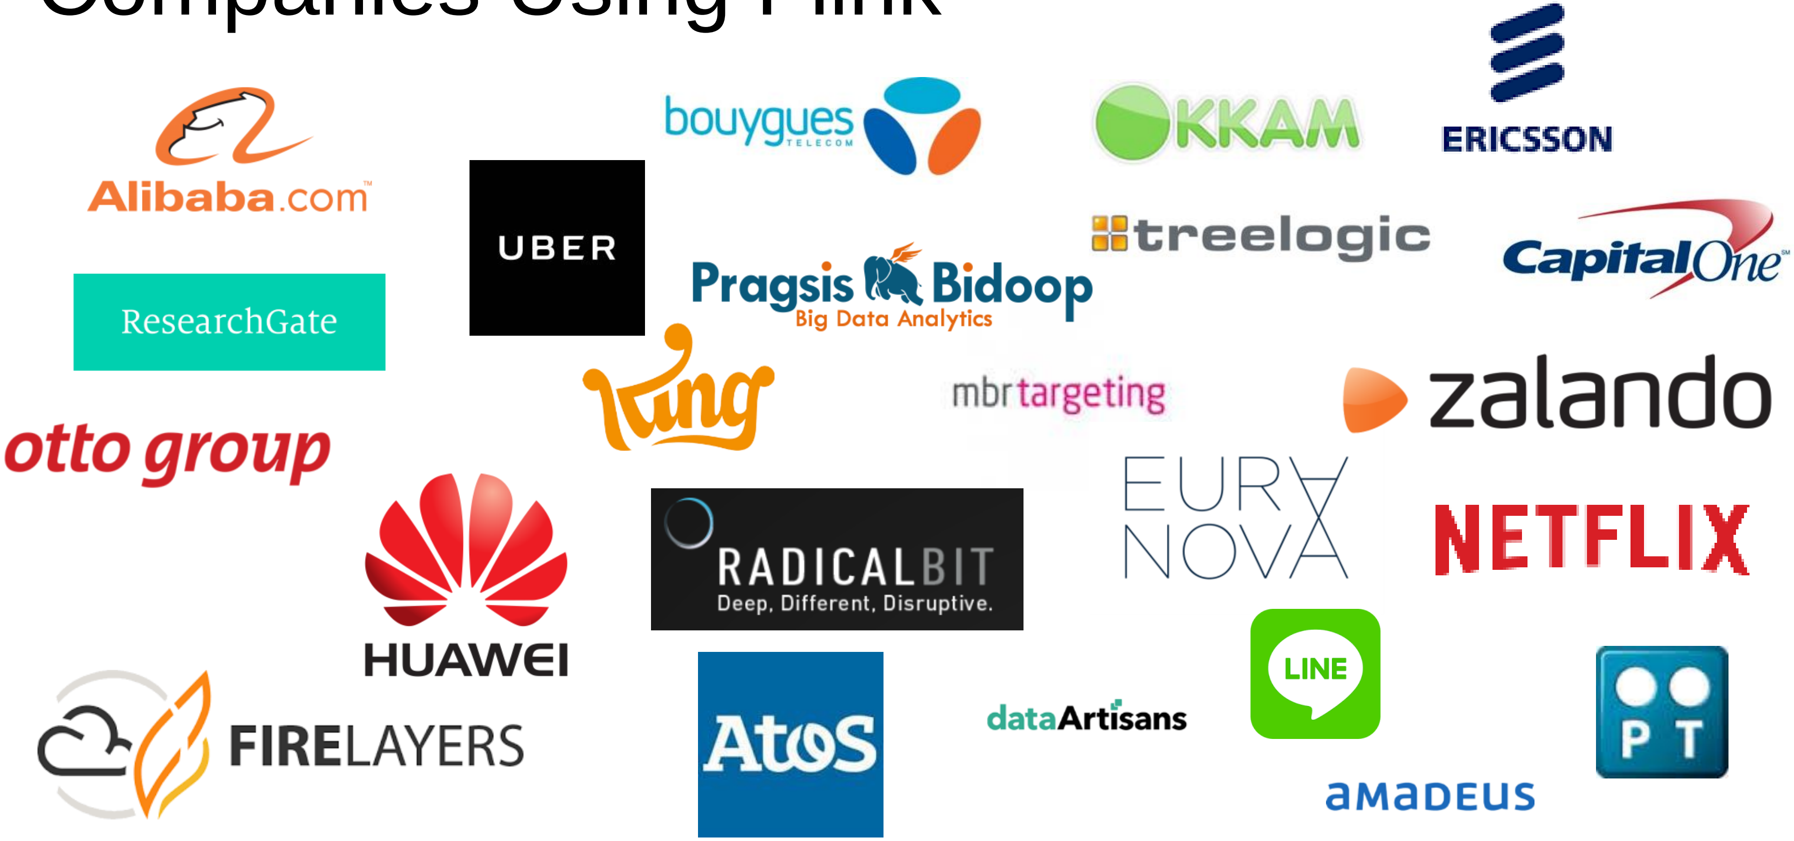

Innovation und Transfer am Beispiel Apache Flink

Flink Community

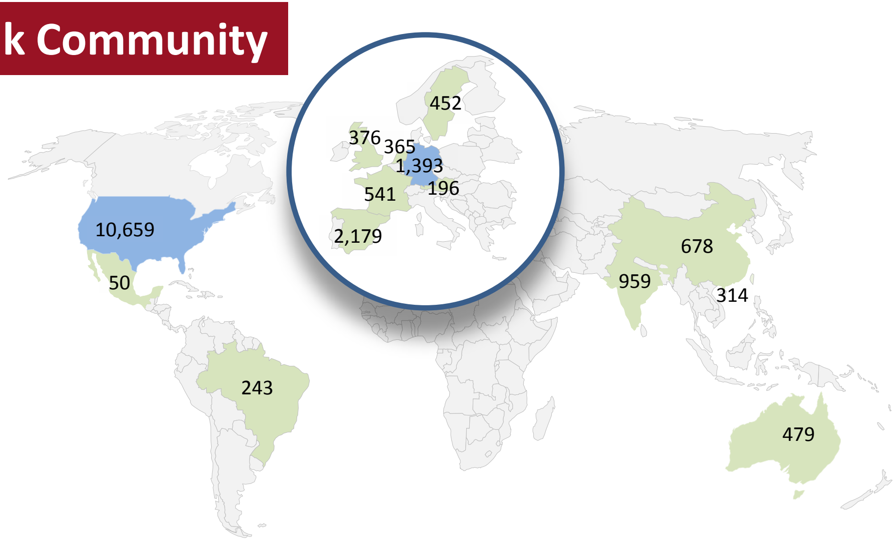

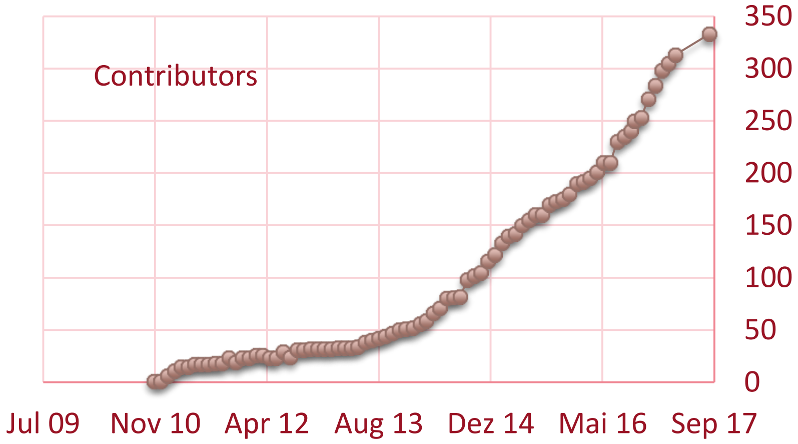

17.820+	Meetup-Mitglieder weltweit	14+	Länder mit regelmäßigen Meetup-Events
328+	open-source Entwickler (contributors)	30+	Anwenderunternehmen
40+	Meetupgruppen weltweit		Firmengründung dataArtisans

Machine Learning + Data Management = X

Mathematical Programming
Linear Algebra
Error Estimation
Active Sampling
Regression Monte Carlo

Relational Algebra/SQL
Data Warehouse/OLAP
NF^2 /XQuery Scalability
Hardware adaption
Fault Tolerance
Resource Management

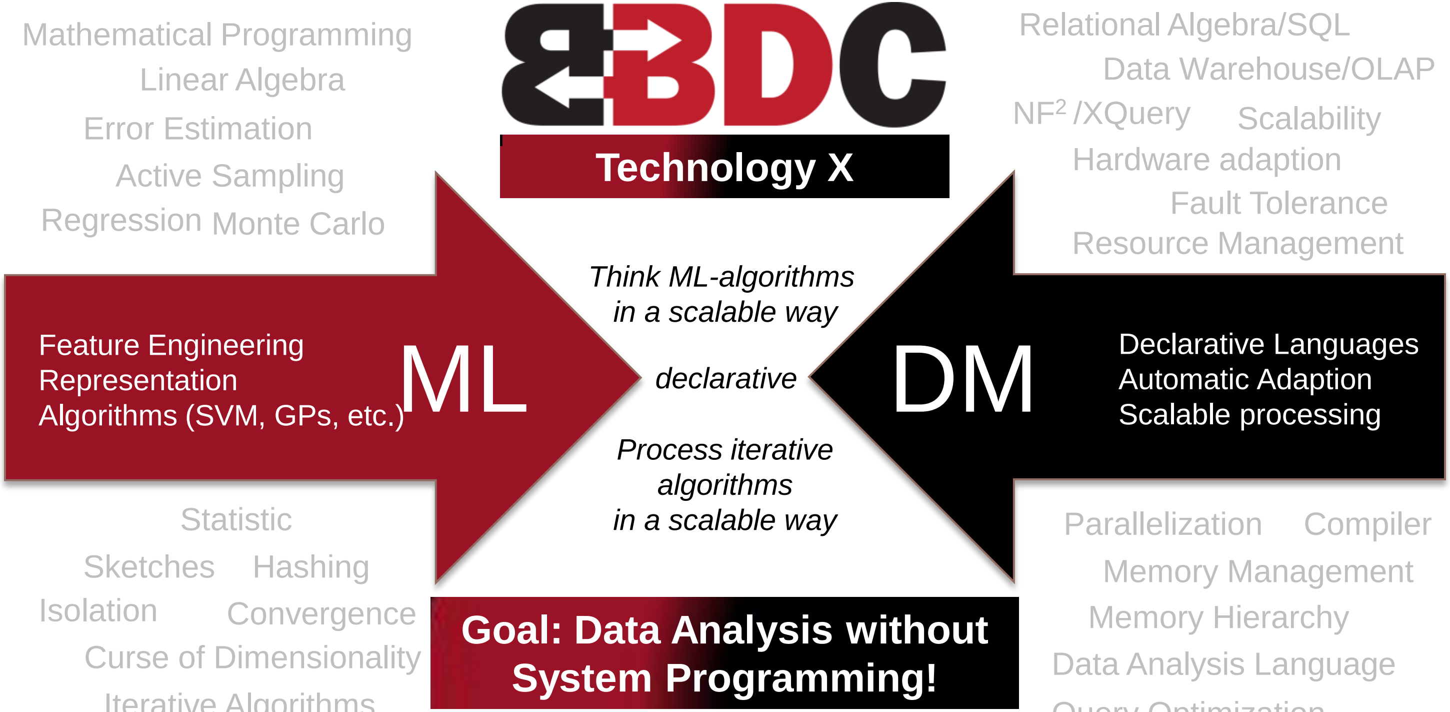

BBDC

Technology X

Feature Engineering
Representation
Algorithms (SVM, GPs, etc.)

ML

Think ML-algorithms in a scalable way

declarative

Process iterative algorithms in a scalable way

DM

Declarative Languages
Automatic Adaption
Scalable processing

Goal: Data Analysis without System Programming!

Statistic
Sketches Hashing
Isolation Convergence
Curse of Dimensionality
Iterative Algorithms
Control flow

Parallelization Compiler
Memory Management
Memory Hierarchy
Data Analysis Language
Query Optimization
Dataflow Indexing

What, Not How! Consider K-means Clustering.

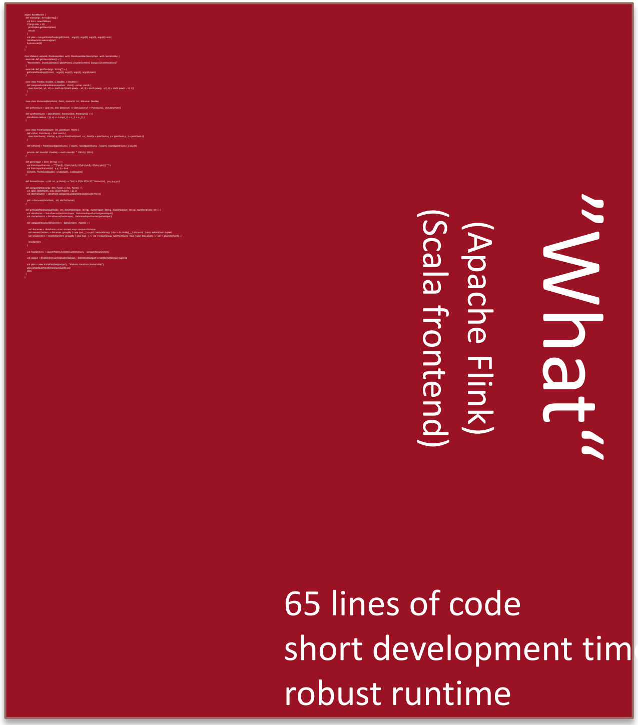
"What"
(Apache Flink)
(Scala frontend)

65 lines of code
short development tim
robust runtime

Declarative data analysis program with automatic optimization, parallelization and hardware adaption

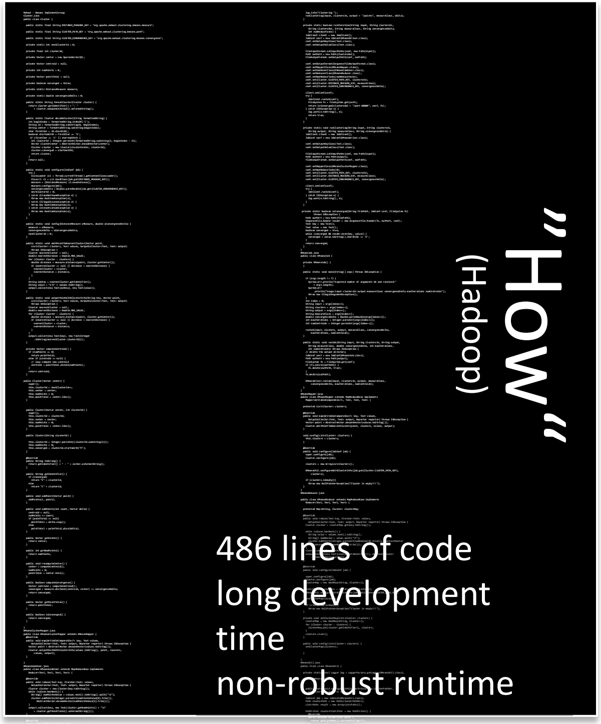

486 lines of code
long development time
non-robust runtime

Hand-optimized code
(data-, load- and system dependent)

Big Data Analytics Without Systems Programming!
(What, Not How!)

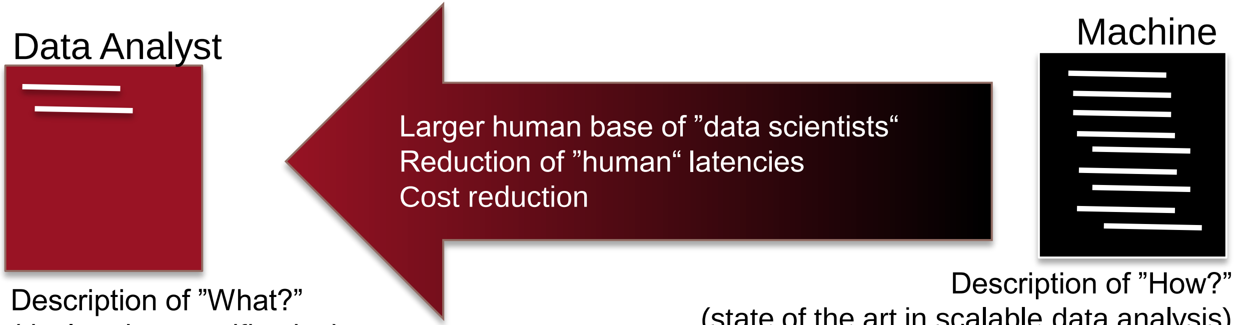

Data Analyst

Larger human base of ”data scientists“
Reduction of ”human“ latencies
Cost reduction

Machine

Description of ”What?”
(declarative specification)

Description of ”How?”
(state of the art in scalable data analysis)
Map/Reduce, MPI

Flink in the BBDC Stack

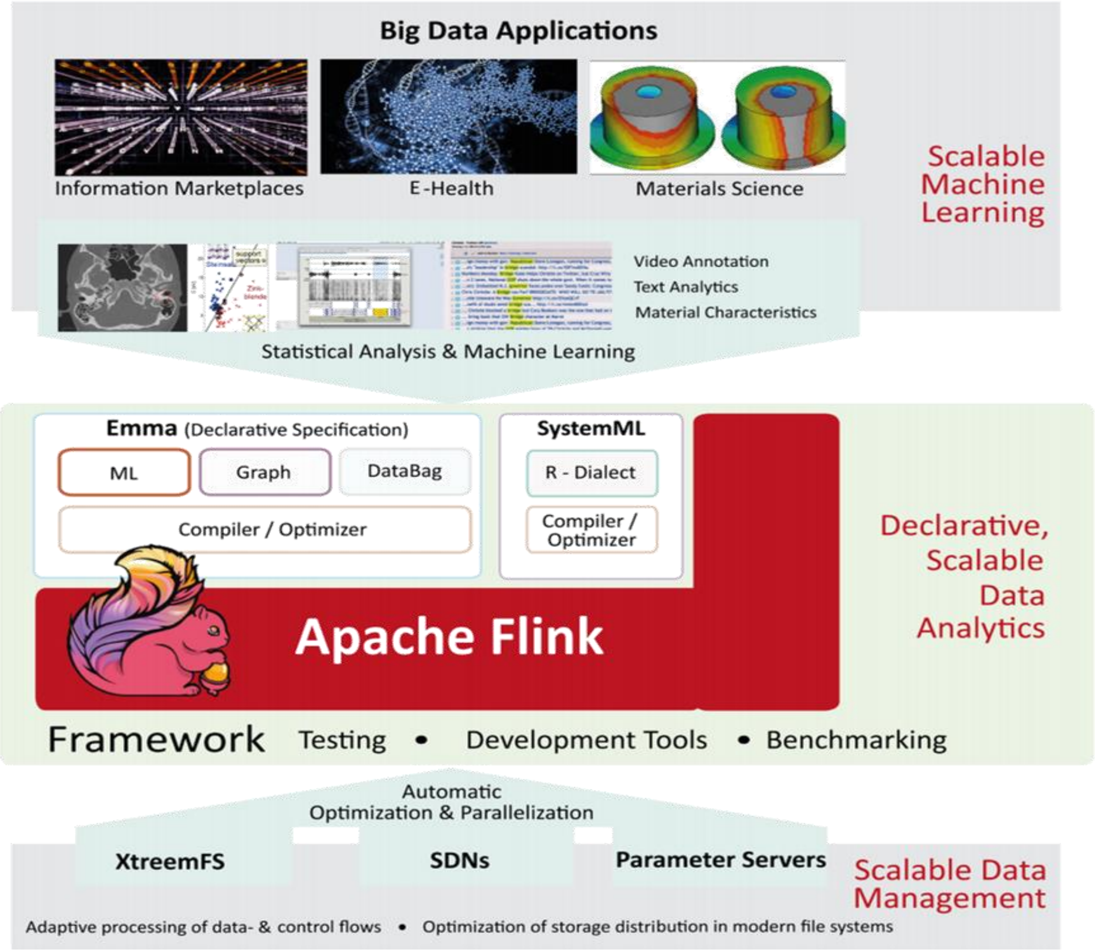

BBDC Lessons Learned & Challenges

- International competitors have far more funding and visibility
  - Berkeley AmpLab: $30 Mio for 6 years, recently grown into Berkeley Institute for Data Science with another $40 Mio funding
  - UK Turing Institute: GBP 67 Mio funding
  - And many others (e.g., across the US, China, Korea, and Japan)

- German companies are followers, not leaders in big data
  - Many of Germany‘s large companies have not yet developed a big data strategy and are risk-averse, or focus too much on short-term and established solutions.
  - It is far easier to work with “new“ companies to transfer novel technologies
    - ResearchGate, Zalando, King, IMR, and Spotify, among others
  - Open source solutions and/or establishing new companies are the best route to turn research into innovation
  - US and large international companies are easier to collaborate with, often times via their respective German subsidiary.

Thank You

Contact:

Jonas Traub

jonas.traub@tu-berlin.de

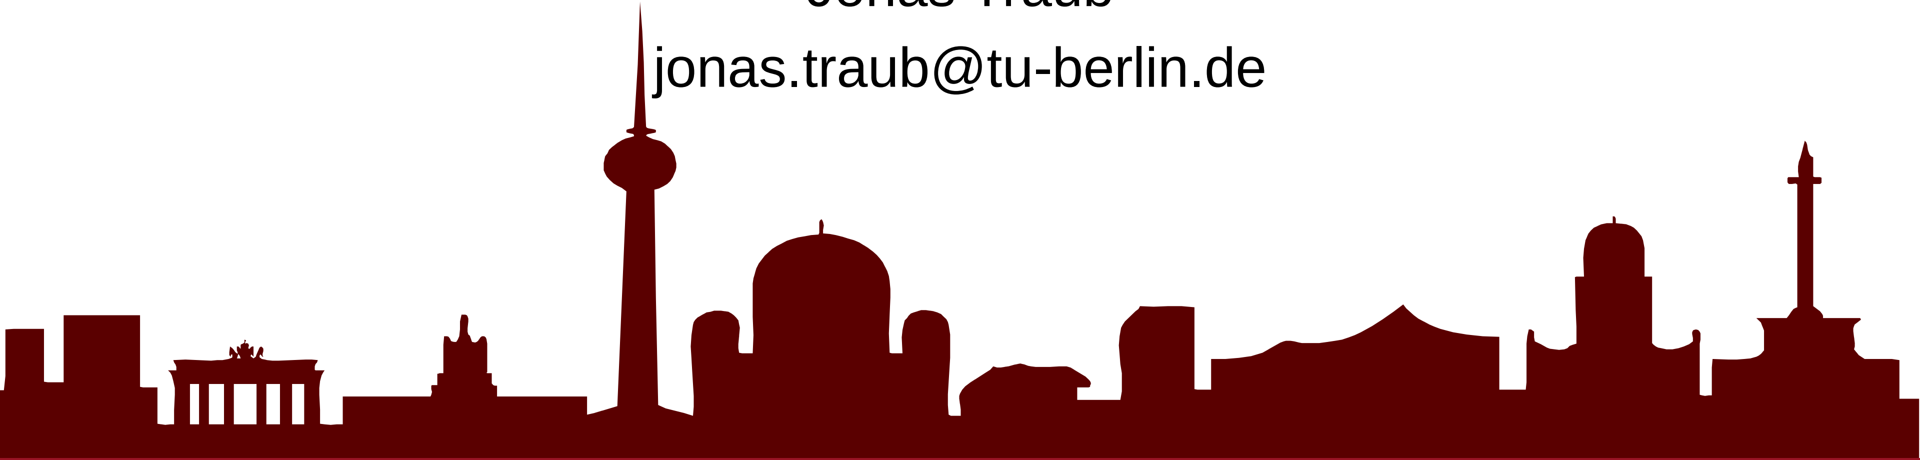

The Berlin Big Data Center

Jonas Traub

Technische Universität Berlin / DFKI IAM

www.dima.tu-berlin.de | bbdc.berlin | jonas.traub@tu-berlin.de

with materials from Tilmann Rabl and Volker Markl