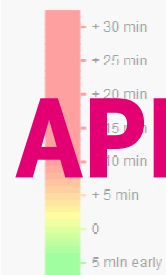


Map by San Jose [GFDL or CC-BY-SA-3.0], via Wikimedia Commons

T-Systems ARRIVAL PROGNOSIS



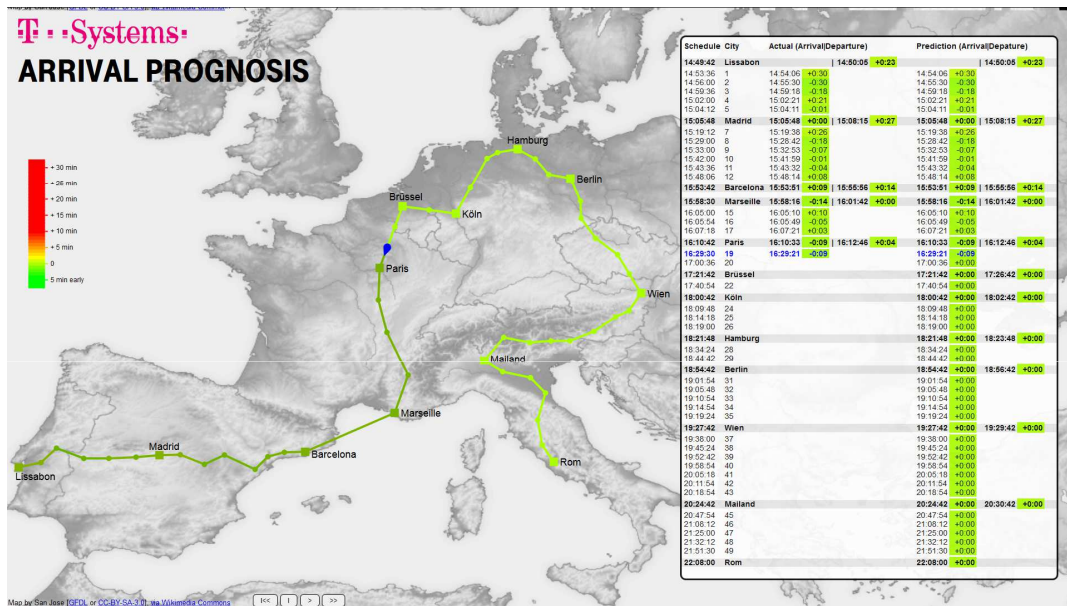
APPLICATION OF MACHINE LEARNING APPROACHES TO REAL-TIME PREDICTION OF TRAIN ARRIVALS

INGO.ELSEN@T-SYSTEMS.COM

Schedule	City	Actual (Arrival Departure)		Prediction (Arrival Departure)	
14:49:42	Lissabon		14:50:05 +0:23		14:50:05 +0:23
14:53:36	1	14:54:06 +0:30		14:54:06 +0:30	
14:56:00	2	14:55:30 -0:30		14:55:30 -0:30	
14:59:36	3	14:59:18 -0:18		14:59:18 -0:18	
15:02:00	4	15:02:21 +0:21		15:02:21 +0:21	
15:04:12	5	15:04:11 -0:01		15:04:11 -0:01	
15:05:48	Madrid	15:05:48 +0:00	15:08:15 +0:27	15:05:48 +0:00	15:08:15 +0:27
15:19:12	7	15:19:38 +0:26		15:19:38 +0:26	
15:29:00	8	15:28:42 -0:18		15:28:42 -0:18	
15:33:00	9	15:32:53 -0:07		15:32:53 -0:07	
15:42:00	10	15:41:59 -0:01		15:41:59 -0:01	
15:43:36	11	15:43:32 -0:04		15:43:32 -0:04	
15:40:06	12	15:40:14 +0:08		15:40:14 +0:08	
16:53:42	Barcelona	16:53:51 +0:09	16:56:58 +0:14	16:53:51 +0:09	16:56:58 +0:14
15:58:30	Marseille	15:58:16 -0:14	16:01:42 +0:00	15:58:16 -0:14	16:01:42 +0:00
16:05:00	15	16:05:10 +0:10		16:05:10 +0:10	
16:05:54	16	16:05:49 -0:05		16:05:49 -0:05	
16:07:18	17	16:07:21 +0:03		16:07:21 +0:03	
16:10:42	Paris	16:10:33 -0:09	16:12:46 +0:04	16:10:33 -0:09	16:12:46 +0:04
16:29:30	19	16:29:21 -0:09		16:29:21 -0:09	
17:00:36	20			17:00:36 +0:00	
17:21:42	Brüssel			17:21:42 +0:00	17:26:42 +0:00
17:40:54	22			17:40:54 +0:00	
18:00:42	Köln			18:00:42 +0:00	18:02:42 +0:00
18:09:48	2			18:09:48 +0:00	
18:14:18	2			18:14:18 +0:00	
18:19:00	2			18:19:00 +0:00	
18:21:48	Hamburg			18:21:48 +0:00	18:23:48 +0:00
18:34:24	28			18:34:24 +0:00	
18:44:42	29			18:44:42 +0:00	
18:54:42	Berlin			18:54:42 +0:00	18:56:42 +0:00
19:01:54	31			19:01:54 +0:00	
19:05:54	32			19:05:54 +0:00	
19:14:44	34			19:14:44 +0:00	
19:19:24	35			19:19:24 +0:00	
19:27:42	Wien			19:27:42 +0:00	19:29:42 +0:00
19:38:00	37			19:38:00 +0:00	
19:45:24	38			19:45:24 +0:00	
19:52:42	39			19:52:42 +0:00	
19:58:54	40			19:58:54 +0:00	
20:05:18	41			20:05:18 +0:00	
20:11:54	42			20:11:54 +0:00	
20:18:54	43			20:18:54 +0:00	
20:24:42	Mailand			20:24:42 +0:00	20:30:42 +0:00
20:47:54	45			20:47:54 +0:00	
21:08:12	46			21:08:12 +0:00	
21:25:00	47			21:25:00 +0:00	
21:32:12	48			21:32:12 +0:00	
21:51:30	49			21:51:30 +0:00	
22:08:00	Rom			22:08:00 +0:00	

Map by San Jose [GFDL or CC-BY-SA-3.0], via Wikimedia Commons

TASK: PREDICT THE ARRIVAL OF TRAINS



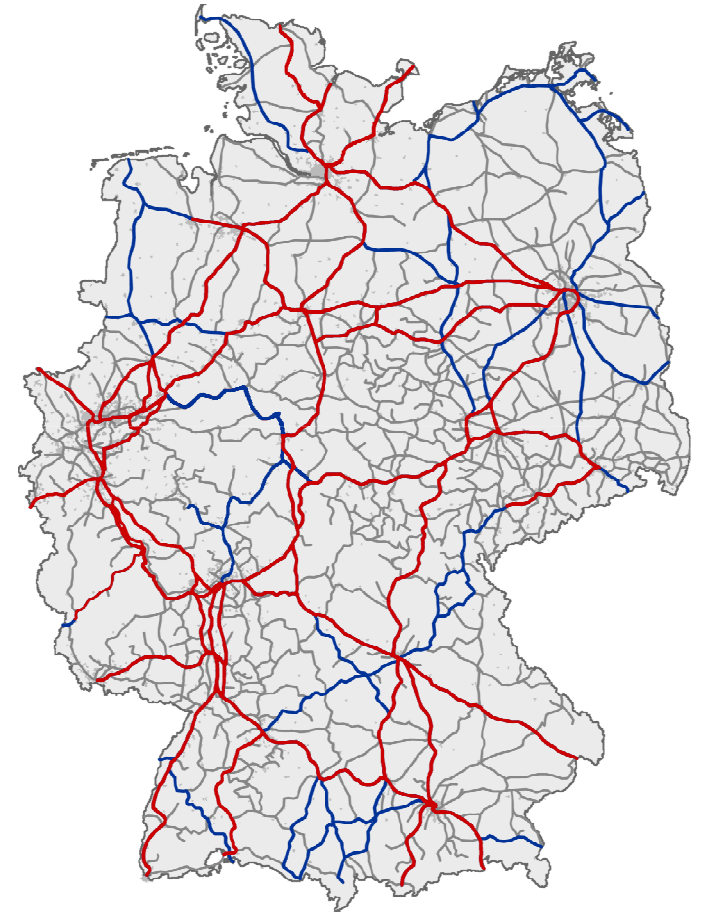
THE FIVE V'S: VOLUME, VELOCITY, VERACITY, VARIETY, VALUE

- 1bn datasets to process for learning
→ 22000 Neural Networks trained in 33 min
- Realtime requirement for operations (sub second latency)
→ <600us to predict a waypoint, two complete operational days predicted with <10min
- Data must be checked for validity
→ validation runs on each message received, prediction quality is continuously checked
- Multiple sources must be merged
→ trains, schedules, network ops systems, ...
- Improve passenger information for running and scheduled trains
→ > two times better than existing solution

MODEL: PREDICT DURATION OF TOUR ALONG A DIRECTED GRAPH

GRAPH PROPERTIES

- Ca. 33000 vertices
- A vertex can be a stop on a tour for one train, but a runthrough point for another
- Some edges in the graph are used by only one type of train
- Most edges are used by multiple types of trains
- Data does not exist for all types (e.g. freight trains are missing but influence overall punctuality)
- Data cannot be expected to be correct at all times



THE AVAILABLE DATA MUST BE PREPROCESSED INTO DIGESTABLE STREAMS

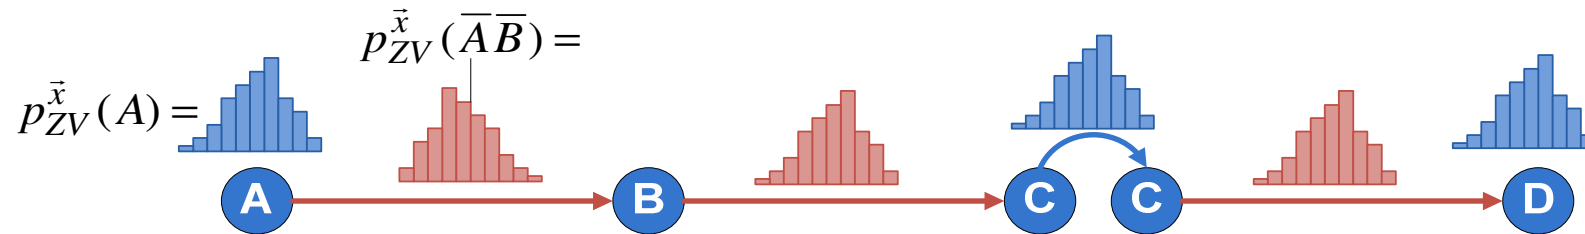
REAL-TIME MESSAGES

f_btg	f_zn	f_nr	f_fsstat	f_bstid	f_rsoll	f_ist	f_sglsid	f_iglsid	f_sstrid	t_nr
21.04.2015	883H	1	2 AK	21.04.2015	14:49	21.04.2015 14:50	103	103	1220	2
21.04.2015	883H	2	5 AMR	21.04.2015	14:53	21.04.2015 14:54	84	84	1220	3
21.04.2015	883H	3	5 AFT	21.04.2015	14:56	21.04.2015 14:56	82	82	1220	4
21.04.2015	883H	4	5 ABRD	21.04.2015	14:59	21.04.2015 14:59	601	601	1220	5

SCHEDULE DATA

lfd_nr	btg	zn	bstid	bstname	halt_lfd_nr	sollan	sollab	sglsid	istan	istab	iglsid	sstrid	unplausibel	vhalt	progan	progab	gattung
1	21.04.2015	883H	AK	Kiel Hbf	0	NULL	21.04.2015 14:49	103	NULL	NULL	103	1220	0	1	False	False	FERN
2	21.04.2015	883H	AMR	NULL	-1	21.04.2015 14:53	21.04.2015 14:53	84	NULL	NULL	84	1220	0	0	False	False	FERN
3	21.04.2015	883H	AFT	NULL	-1	21.04.2015 14:56	21.04.2015 14:56	82	NULL	NULL	82	1220	0	0	False	False	FERN
4	21.04.2015	883H	ABRD	NULL	-1	21.04.2015 14:59	21.04.2015 14:59	601	NULL	NULL	601	1220	0	0	False	False	FERN
5	21.04.2015	883H	AEF	NULL	-1	21.04.2015 15:02	21.04.2015 15:02	25BP	NULL	NULL	25BP	1220	0	0	False	False	FERN
6	21.04.2015	883H	AN G	NULL	-1	21.04.2015 15:04	21.04.2015 15:04	14	NULL	NULL	14	1220	0	0	False	False	FERN
7	21.04.2015	883H	AN	Neumünster	1	21.04.2015 15:05	21.04.2015 15:07	104	NULL	NULL	104	1220	0	1	False	False	FERN
8	21.04.2015	883H	ABRS	NULL	-1	21.04.2015 15:15	21.04.2015 15:15	202	NULL	NULL	202	1220	0	0	False	False	FERN
9	21.04.2015	883H	AWST	NULL	-1	21.04.2015 15:19	21.04.2015 15:19	2	NULL	NULL	2	1220	0	0	False	False	FERN

CORE CONCEPT: CUMMULATION OF PREDICTED ADDITIONAL DELAYS ALONG THE TRAIN TOUR



Expectation value of distributions

$$t_{pr}(D) = t_{akt} + E[p_{ZV}^{\vec{x}}(\Delta t_{AB})] + E[p_{ZV}^{\vec{x}}(\Delta t_{BC})] + E[p_{ZV}^{\vec{x}}(\Delta t_{CC})] + E[p_{ZV}^{\vec{x}}(\Delta t_{CD})]$$

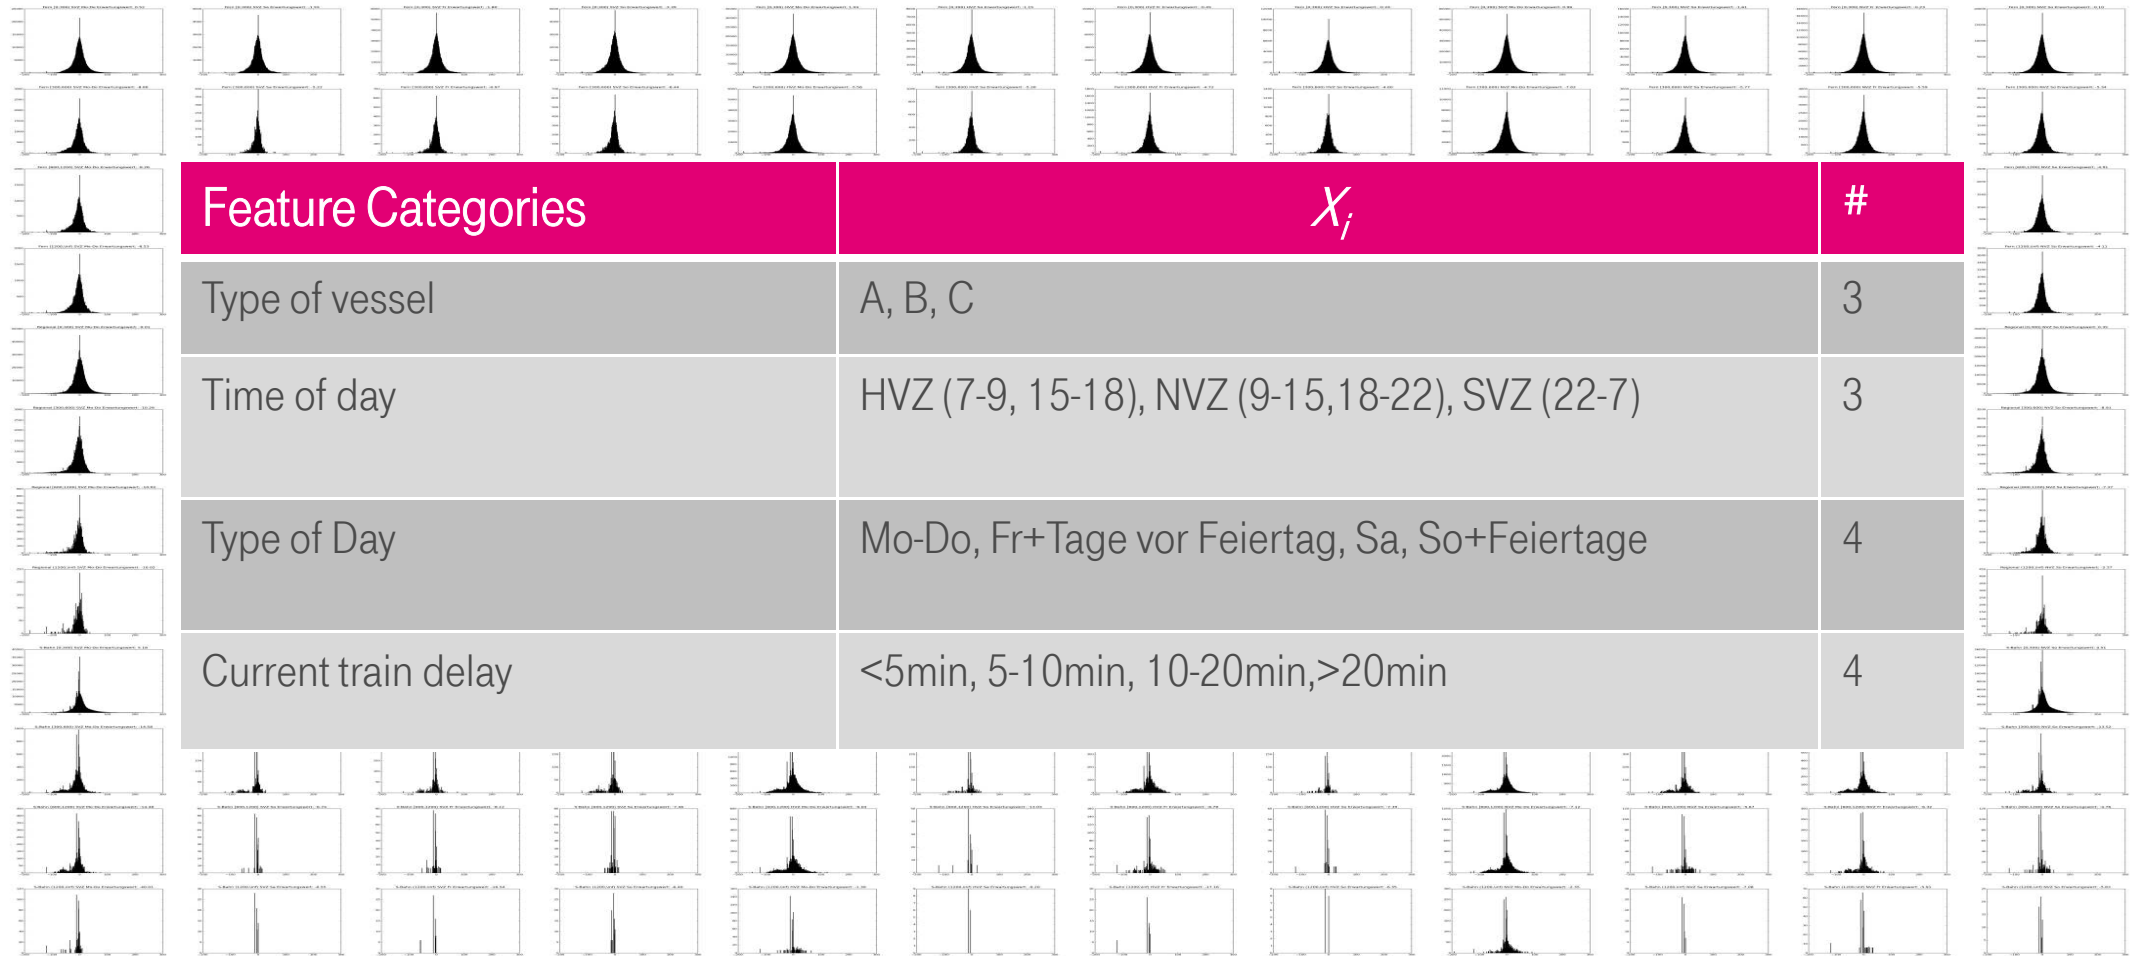
Expectation value of convoluted probability density functions

$$t_{pr}(D) = t_{akt} + E\left[\left(\left(p_{ZV}^{\vec{x}}(\Delta t_{AB}) * p_{ZV}^{\vec{x}}(\Delta t_{BC})\right) * p_{ZV}^{\vec{x}}(\Delta t_{CC})\right) * p_{ZV}^{\vec{x}}(\Delta t_{CD})\right]$$

Sum of local predictions of artificial neural network

$$t_{pr}(D) = t_{akt} + f_{AB}(\vec{x}) + f_{BC}(\vec{x}) + f_{CC}(\vec{x}) + f_{CD}(\vec{x})$$

MODEL FOR NON LEARNING APPROACH: VERTEX SPECIFIC HISTOGRAMS



THE PROBLEMS OF THE „CLASSICAL“ APPROACHES LEAD TO WEAK PREDICTIONS

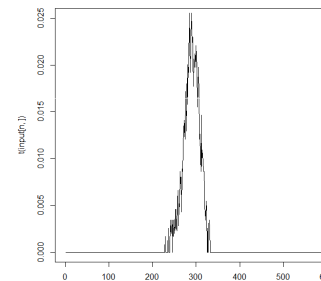
TODAY'S SOLUTION

- Linear extrapolation of current delay spiced with known events (sometimes)
- Data is not always accurate due to manual intervention
- Completely agnostic w.r.t. already known information, e.g. type of vessel, time of day, ...

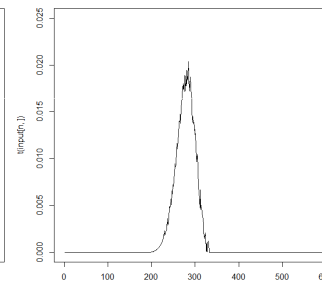
CONVOLUTION OF PDF'S

- For some segments the PDF's are not very “smooth”
- Even for smaller tours, the PDF's for tour segments far in the future are very wide => Low p for point estimation

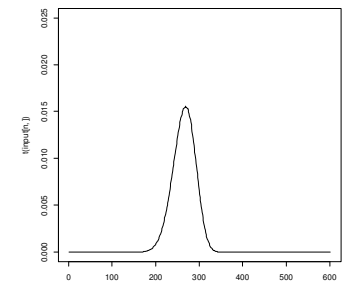
$$(f * g)(n) = \sum_{m=-M}^M f(n-m)g(m)$$



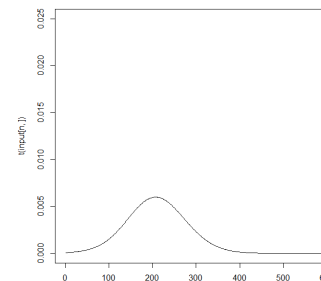
n = 1



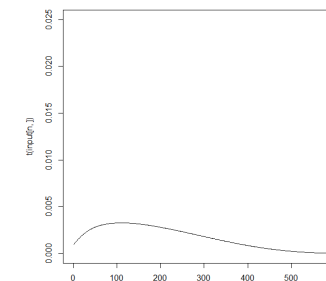
n=2



n = 3

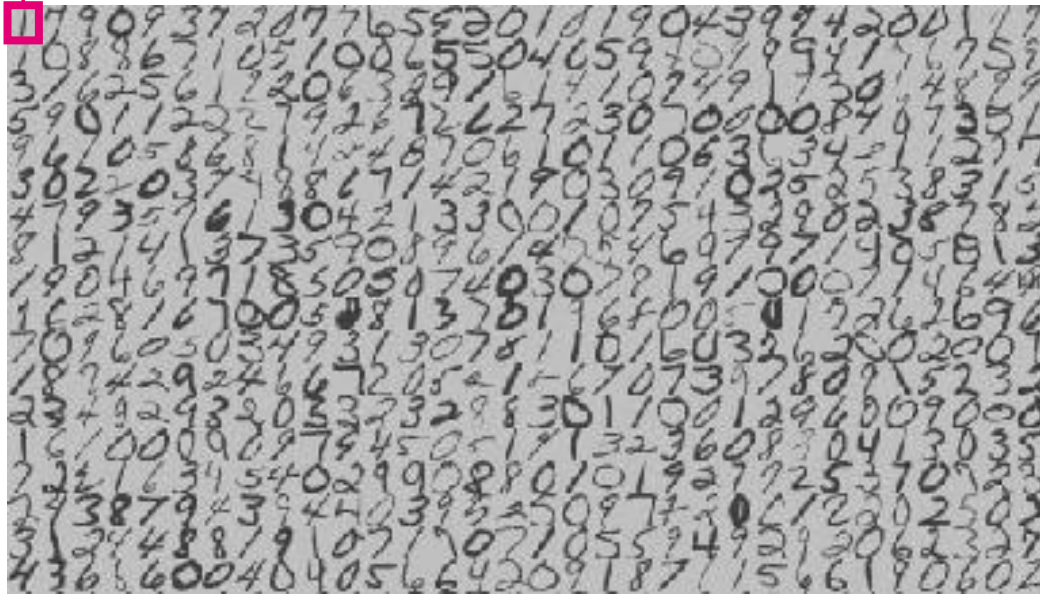


n = 20



n = 150

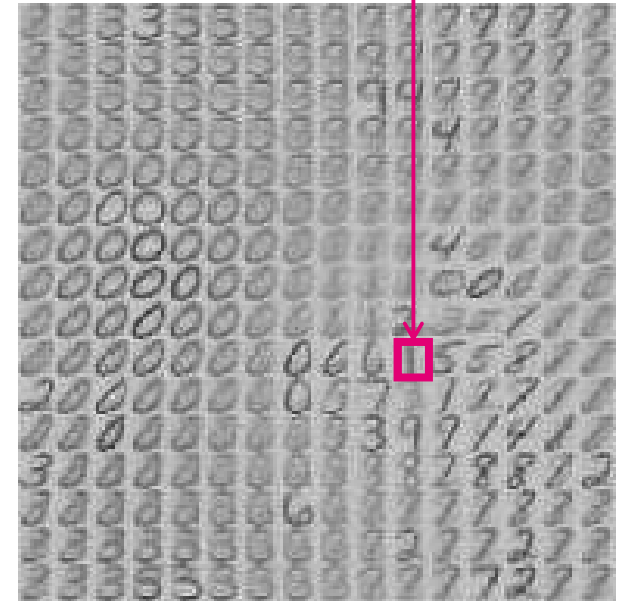
VISUALIZING CLUSTERS OF INFORMATION WITH SELF-ORGANIZING FEATURE MAPS



Input Space

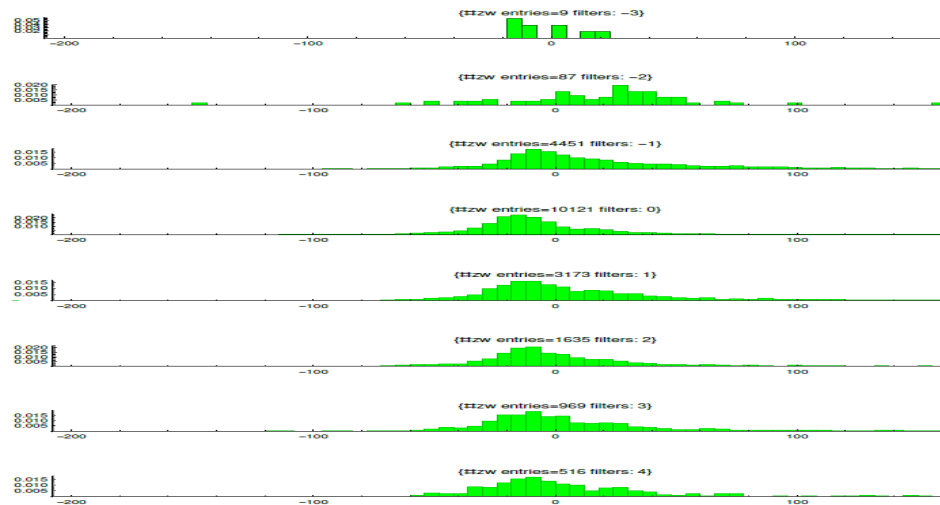
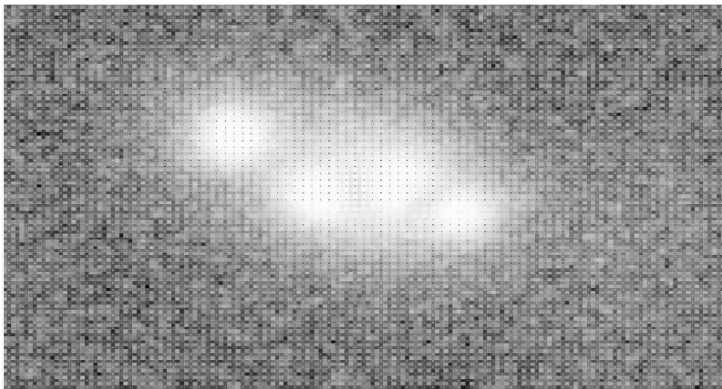
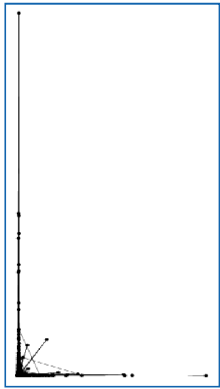
SOFMs are

- dimension-reducing
- piecewise connected
- nonlinear-mapping
- topology-preserving
- variable resolution



Output Space

CLUSTERING SHOWS THAT INCREASE OF DELAY IS NOT CLEARLY SEPARABLE BETWEEN DIFFERENT CLUSTERS



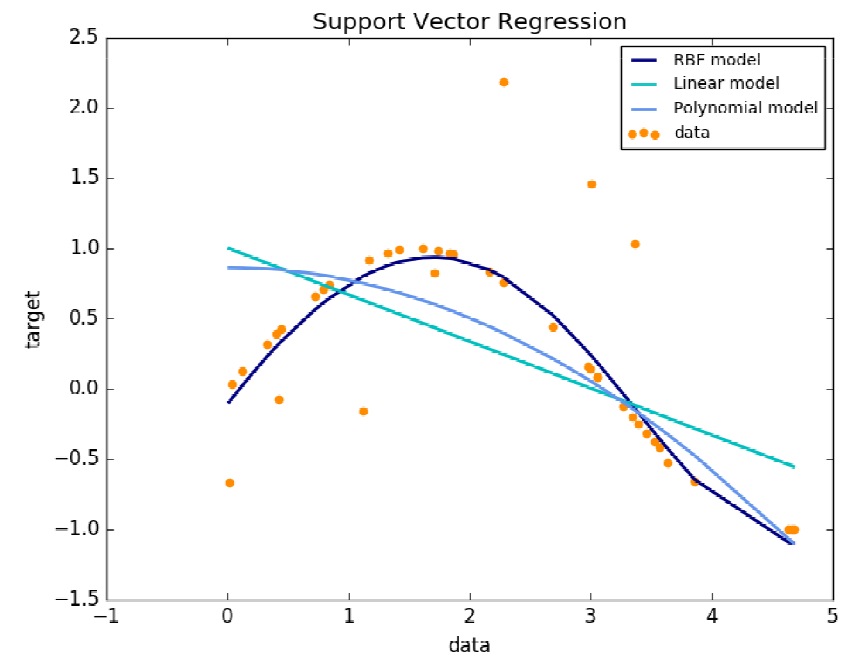
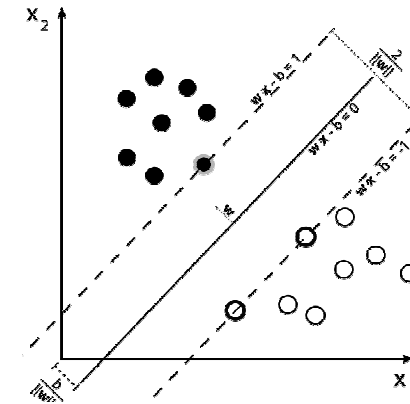
USING SUPERVISED AND NON-SUPERVISED ALGORITHMS

- Optimization of feature categories with unsupervised learning (Self-Organizing Map, k-Means) on different vertices
 - No clear clustering (in terms of feature to detect) with currently available features
 - No usable predictions
- Supervised learning using the existing features provides much better predictions
 - Support Vector Machines (SVM)

SUPERVISED LEARNING APPROACH: NON-LINEAR REGRESSION USING A SUPPORT VECTOR MACHINE

SVM PROPERTIES

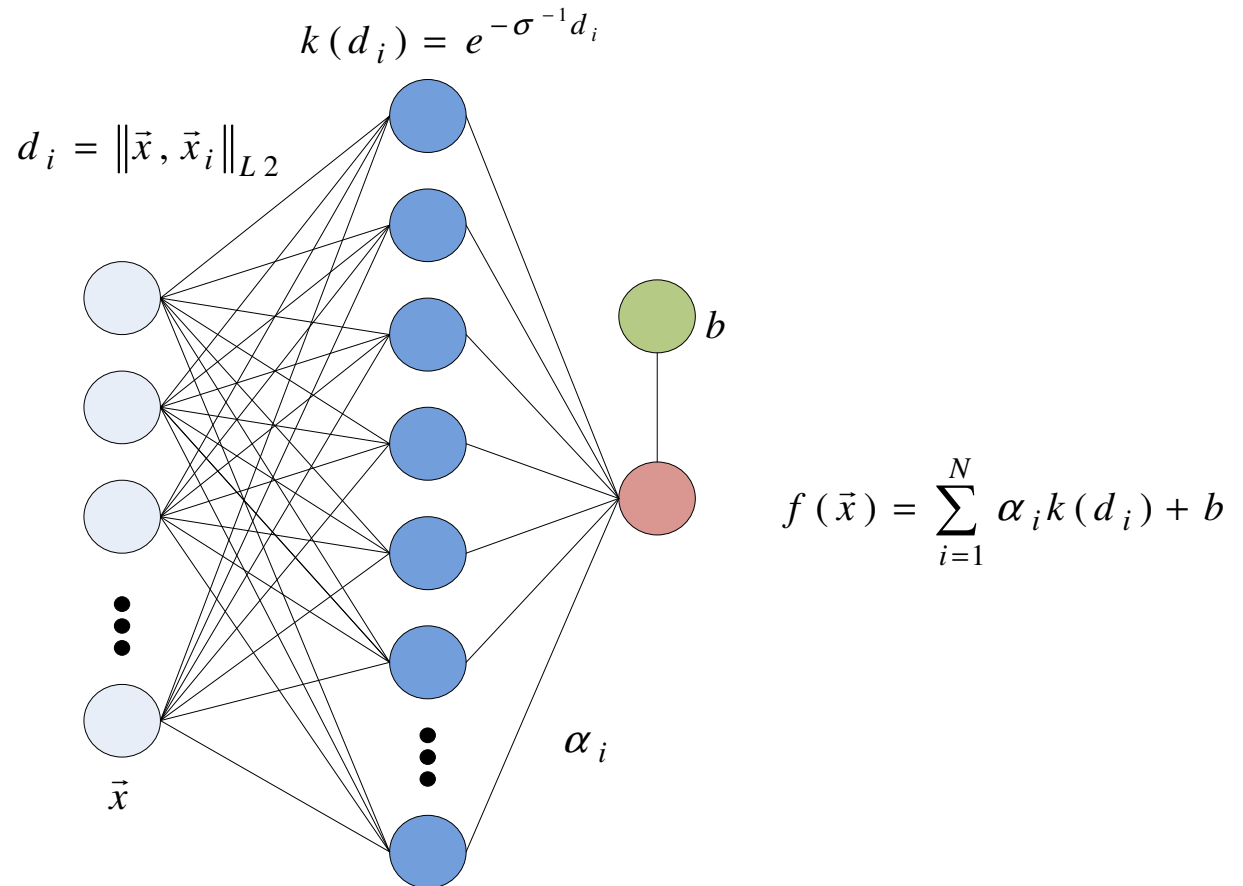
- Finds optimal separating hyperplane in high-dimensional space
- Automatically prunes the number of support vectors needed
- Uses “kernel trick” to make sets separable by transforming the data to higher dimension
- Training is basically an quadratic optimization problem
- svms have found widespread use in classification and prediction



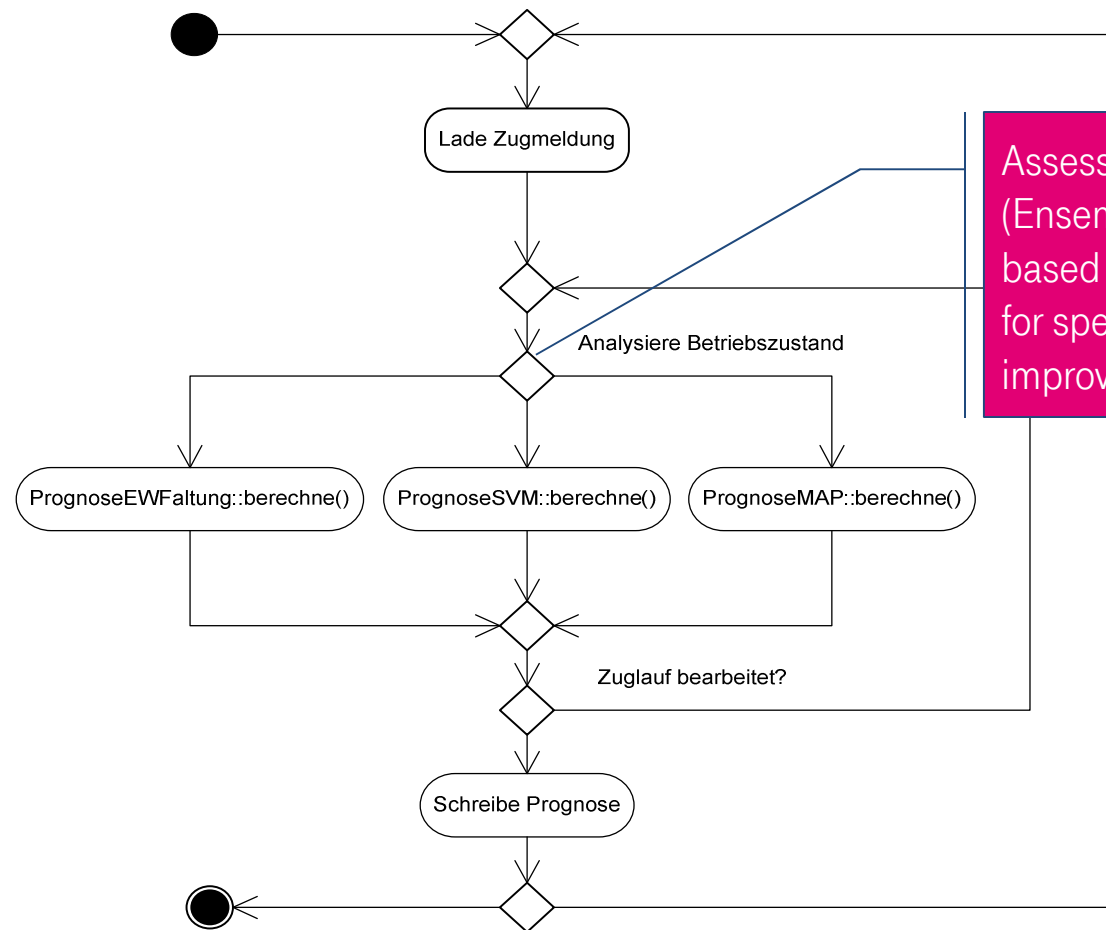
THE NEURAL NETWORKS ARE TRAINED WITH A VARIABLE SET OF INPUT PARAMETERS IN A HIGHLY PARALLELIZED APPROACH

TRAINING OF SVM

- Vertex specific training, i.e. one aNN per vertex
- Highly parallelizable on existing cluster – due to independence of vertices
- Performance: 22000 aNNs in 33 min
- Target value: additional delay



THE DESIGN OF THE SOFTWARE ENABLES THE ONLINE SELECTION OF THE BEST PERFORMING ALGORITHM - ENSEMBLING



Assessment of operational mode:
(Ensemble Method): dynamic selection
based on the performance of the algorithms
for specific operational states (continuously
improved)

IMPROVEMENT OF PREDICTION BY WEIGHTING OF HISTORICAL DATA

YOUNGER DATA IS MORE IMPORTANT

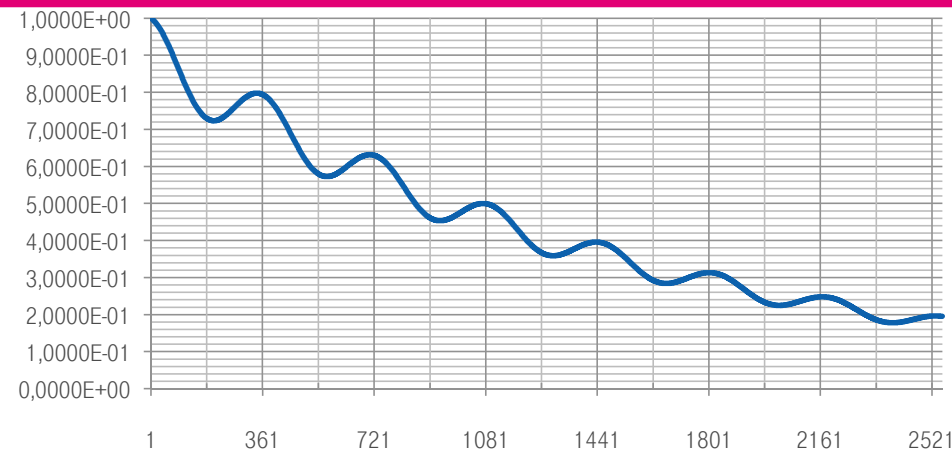
- Higher weight on more current data, e.g. construction work
- Reduction of influence of older situations
- Consideration of seasonality („last july is more important than last september“)

REALIZATION

- Standard approach: exponential filtering
 - Does not take seasonality into account
- Approach: modulate exponential filtering with periodical function

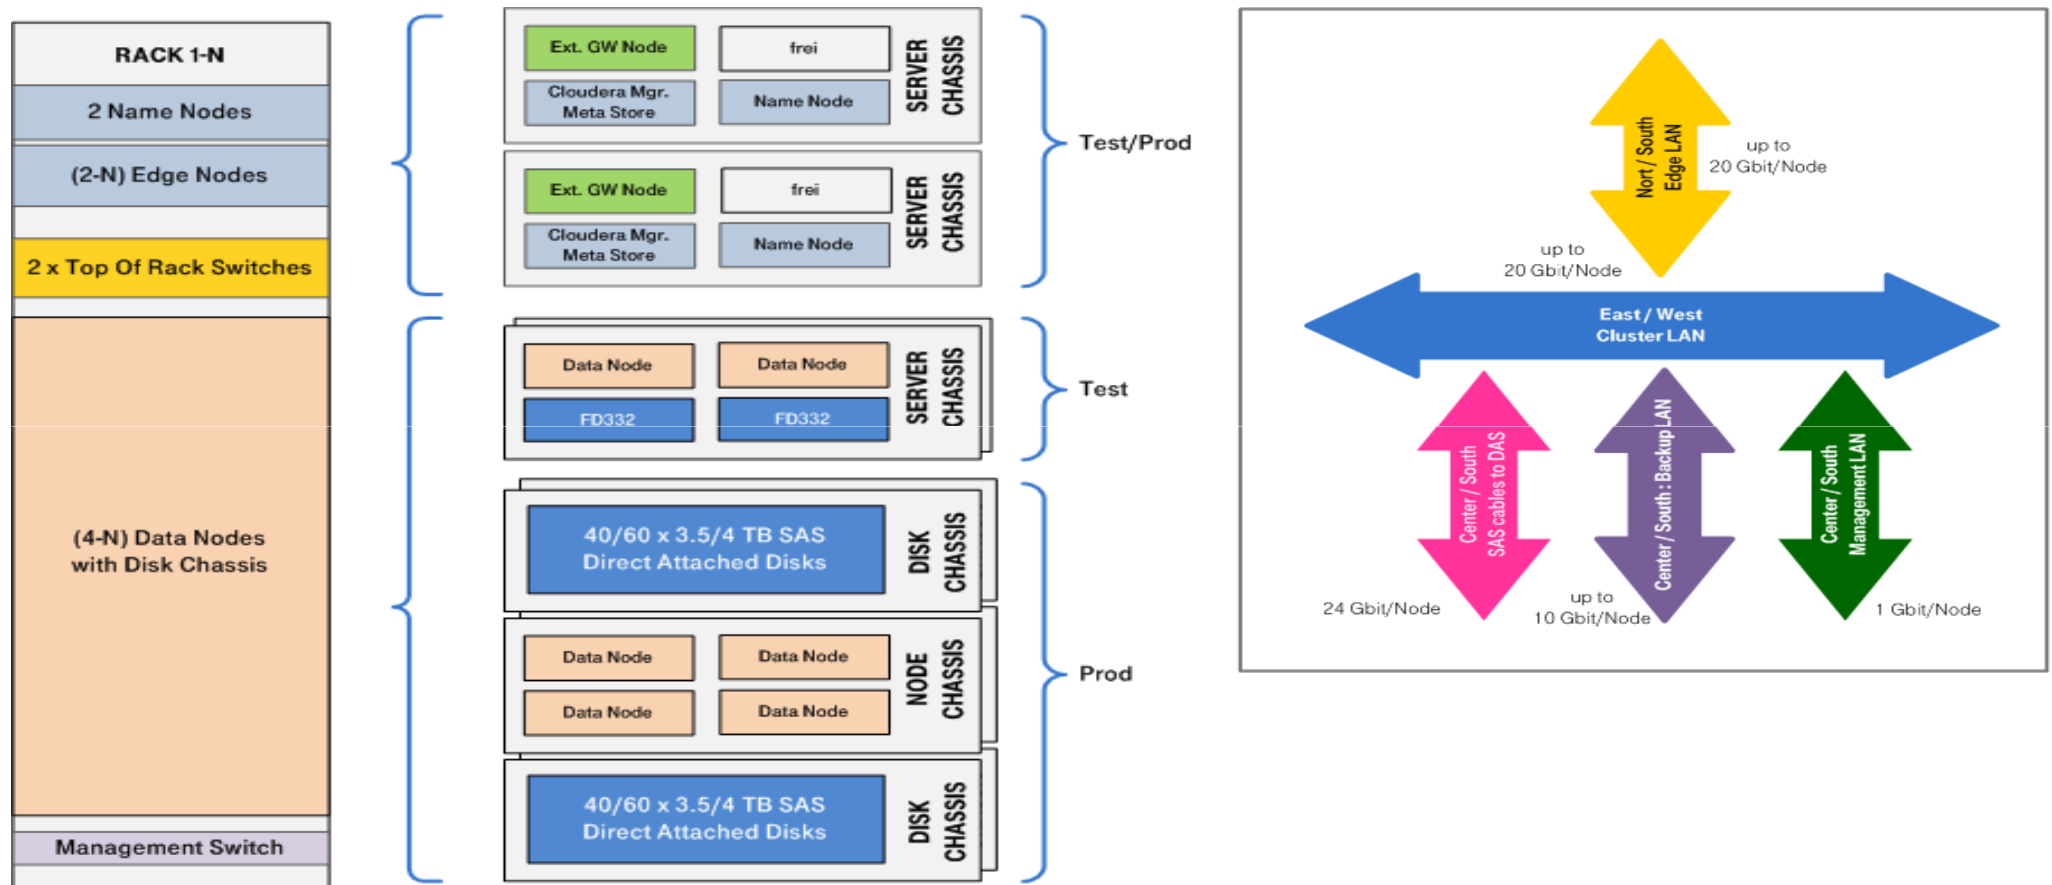
$$f(t) = \alpha \cdot (1 - \alpha)^n \cdot (\eta + \cos(\frac{n}{182.5} \cdot \pi)) \cdot f(0)^{-1}$$

SEASONALIZATION WITH EXPONENTIAL FILTERING



ARCHITECTURE

DATA PLATFORM SYSTEM ARCHITECTURE

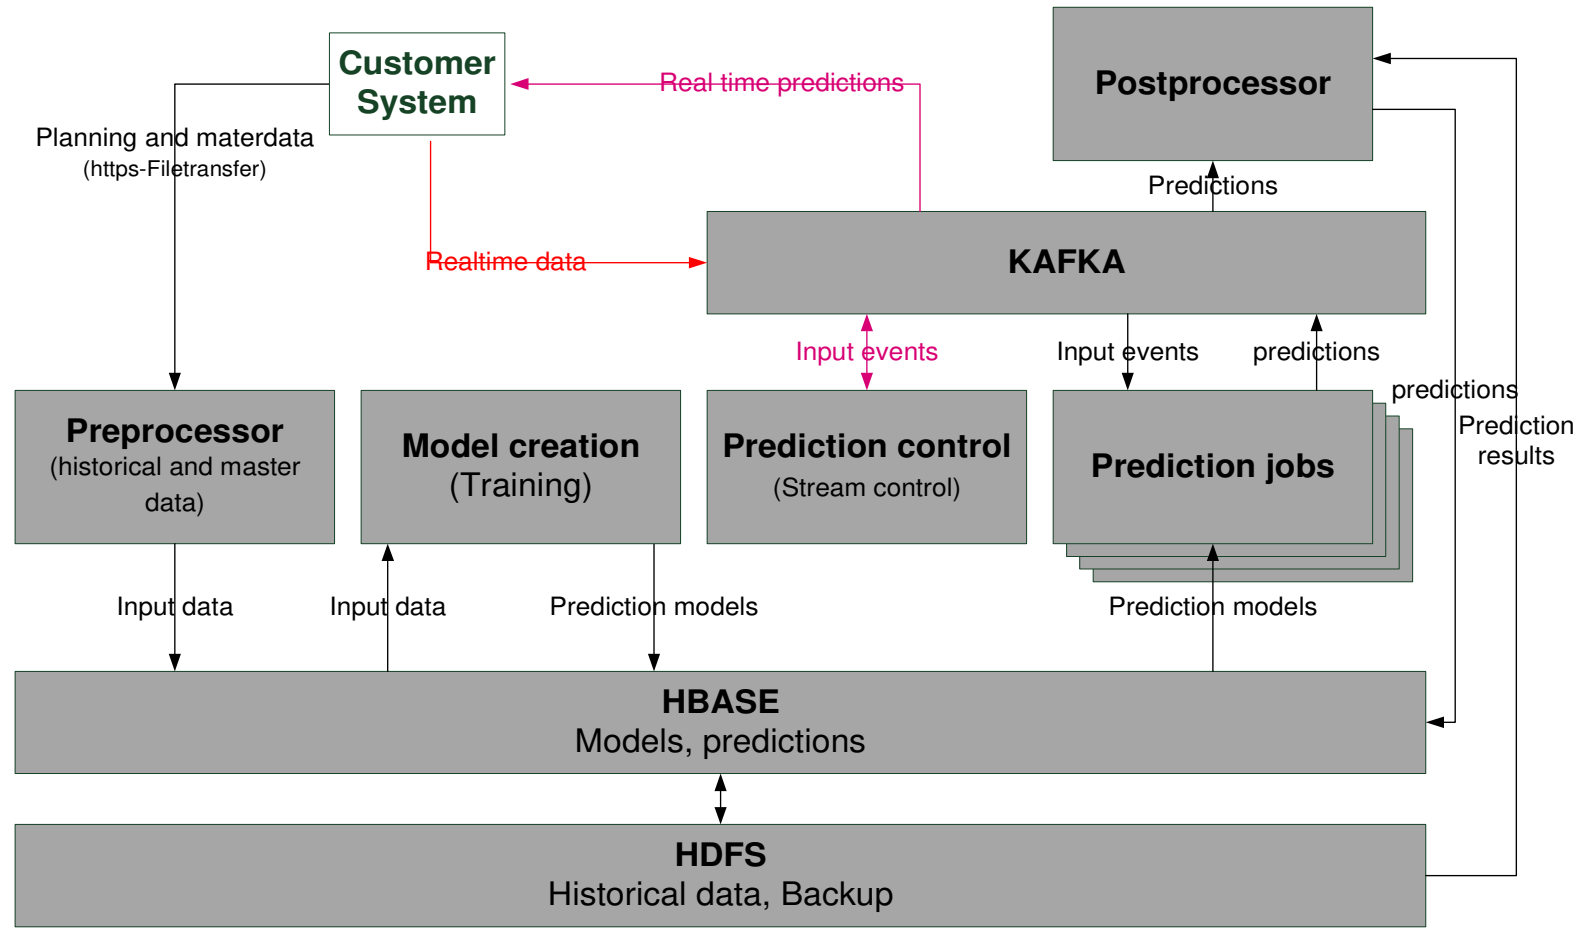


OUR PLATFORM SCALES FROM 3 TO 2400 DATA NODES

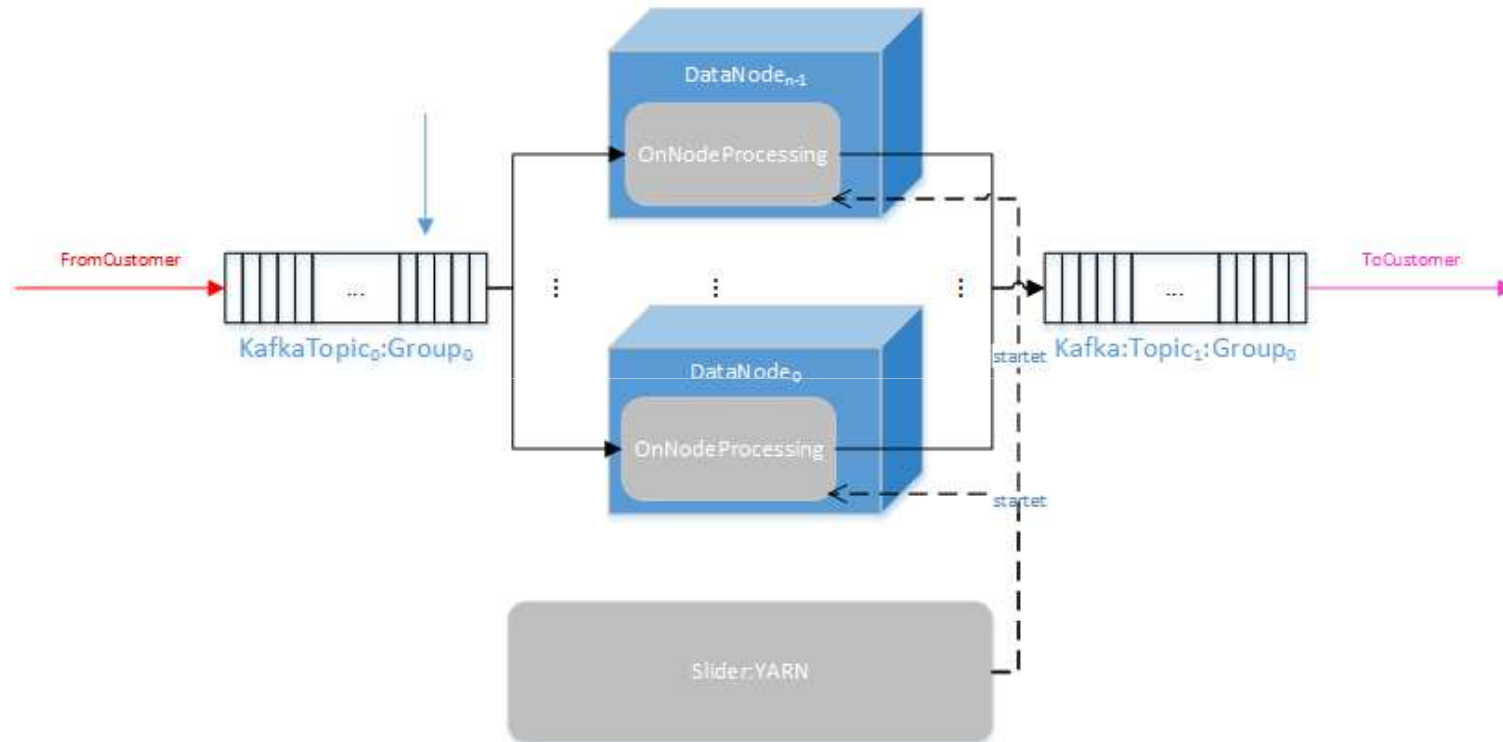
Function	Master & Name Node	Kafka (I/O) Node	Data Node
Hardware	Intel x86		
OS	Linux 64 bit		
CPU	2x Intel Xeon E5-2620 v4 (8 Cores, 20MB Cache)		
RAM	256 GB		
Storage	2 x 2 TB		15 x 4 (8) TB

- 2x10 Gbit/s ToR switches in fabric mode
- >20Gbit/s I/O bandwidth/data node
- Scalable per Rack up to 50 nodes

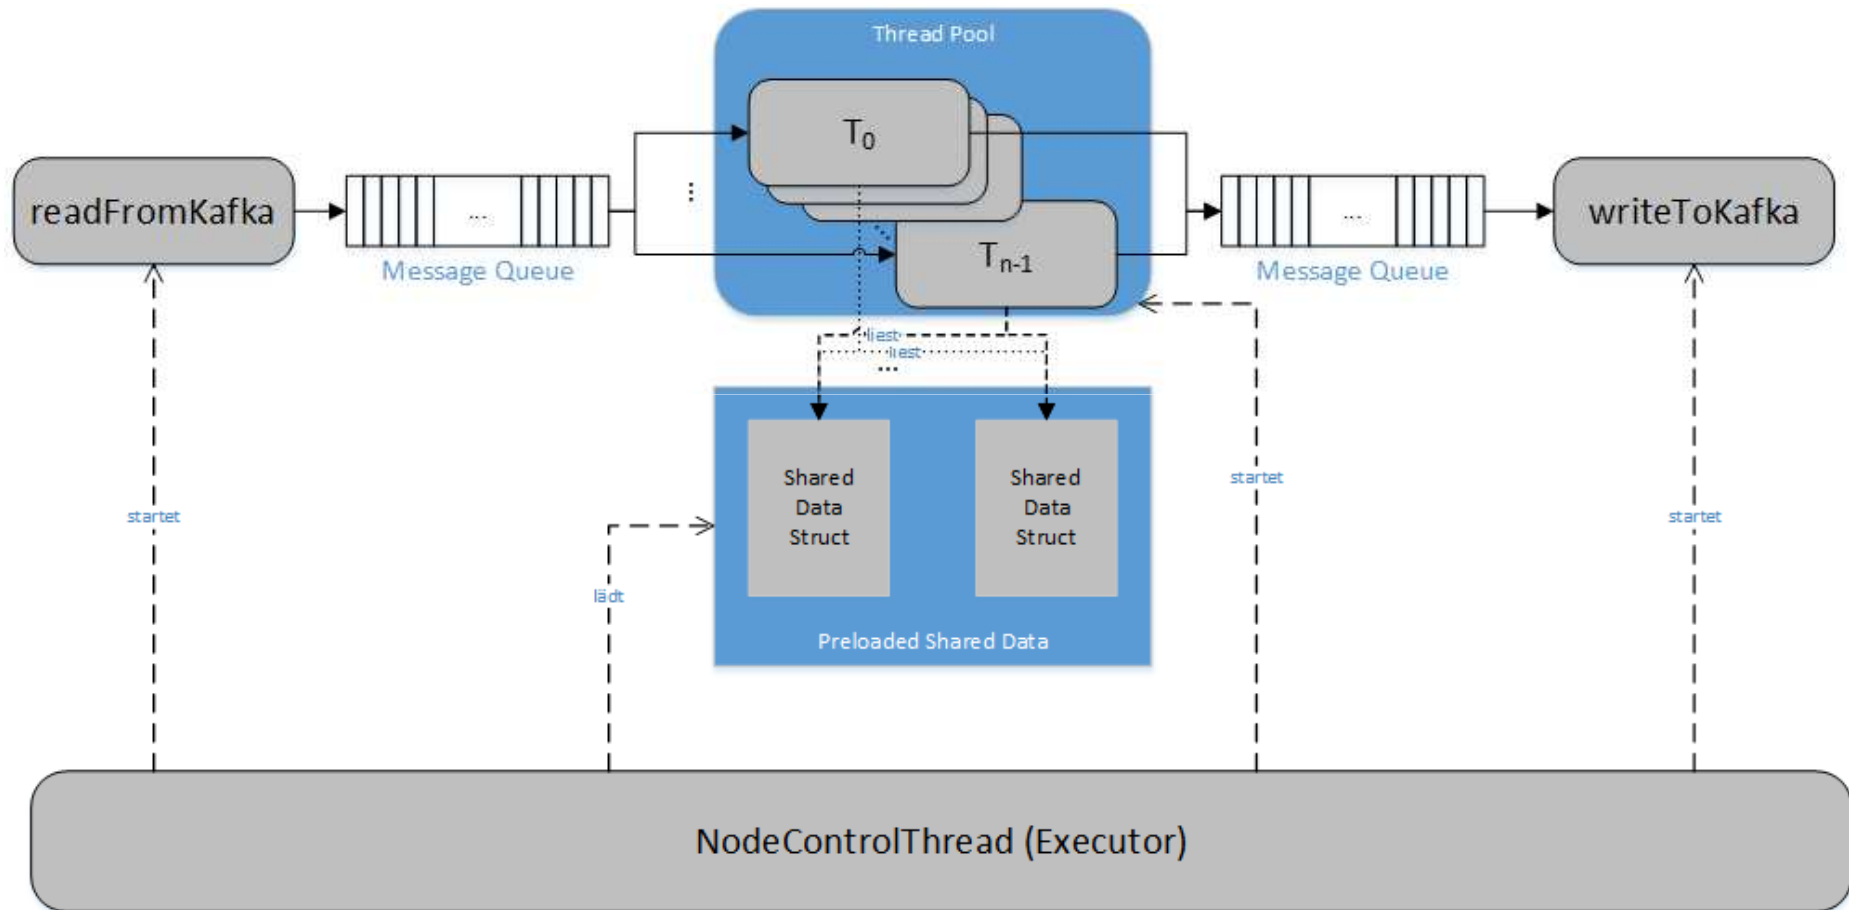
THE TARGET ARCHITECTURE CAN BE FLEXIBLY USED FOR DIFFERENT USE CASES



THE ARCHITECTURE USES THE MINIMAL POSSIBLE SET OF FRAMEWORKS TO IMPLEMENT A FARMER WORKER PATTERN

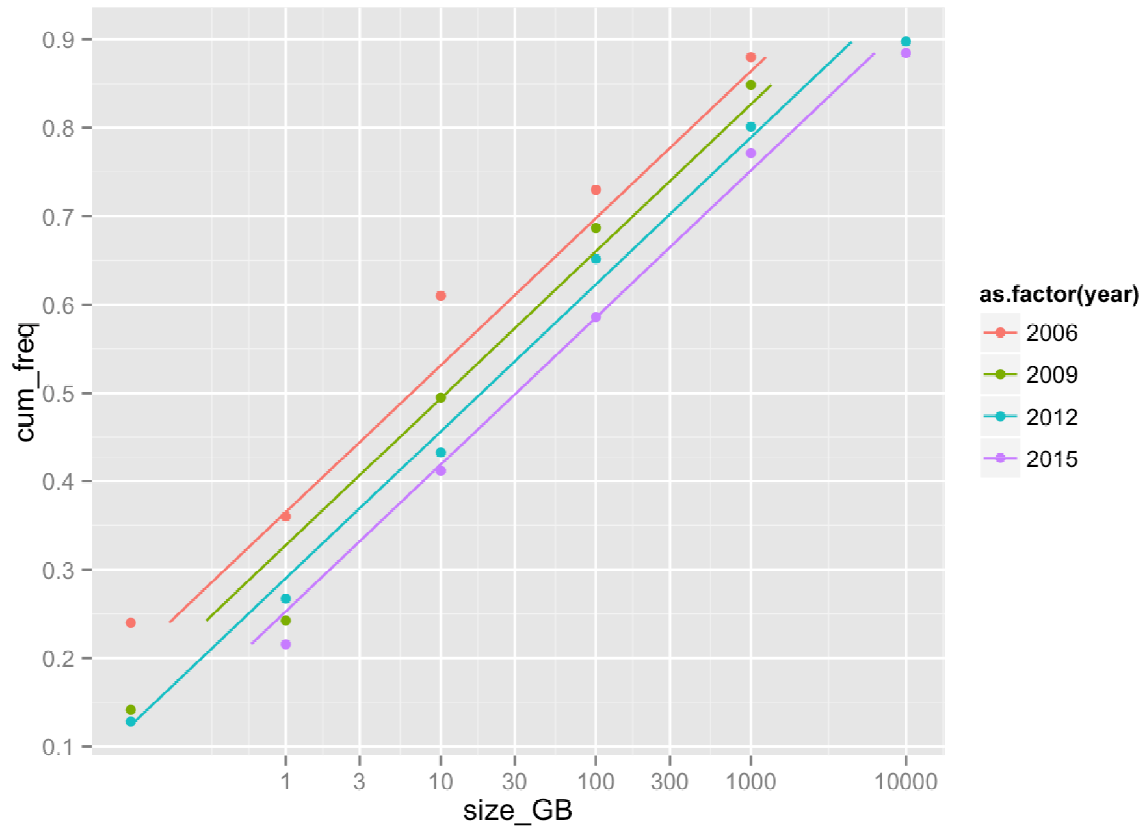


**EACH NODE REPEATS THE FARMER WORKER PATTERN.
IN-MEMORY COMPUTING SPEEDS UP PERFORMANCE.**



TRENDS IN DATA ANALYTICS AND MACHINE LEARNING FROM A COMPUTING PERSPECTIVE

FAST DATA EATS BIG DATA



source: <https://github.com/szilard/dataset-sizes-kdnuggets>

GROWTH RATE IN RAM EXCEEDS GROWTH RATE IN DATA

- Data Node C4 (2013: 32 GiB)
- Data Node DB (2016: 256 GiB)
- RAM Growth/a: 71%
- Data Growth/a: 20%

REAL-TIME PROCESSING BECOMES REALITY

NEW (OLD) DATATYPES ENABLE IN-MEMORY COMPUTING – FP16 (IEEE 754-2008 STANDARDIZED)

FP16 PROPERTIES

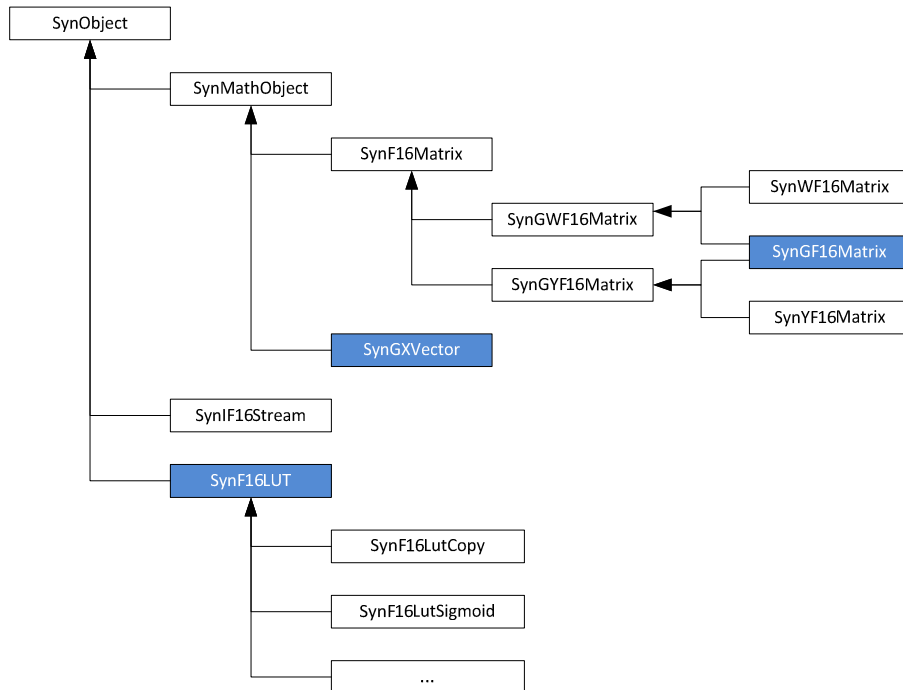
- 16 bit “binary16”
- 10 bit mantissa, 5 bit exponent (+/- 65504.0)
- Originally only intended for storage, not for computation
- Support in x86 since IvyBridge “`__m128 _mm_cvtph_ps (__m128i a)`”
- Now computations are also available in GPU (since nVidia „Pascal“)
- OpenCL as flexible programming framework
- And people tend repeat the errors from the past (SysArr) (arXiv:1502.02551v1 [cs.LG] 9 Feb 2015)

BUT SIMILAR APPROACHES EXISTED ALREADY 25 YEARS AGO

- Most Neural Networks are fine with 16bit floats! (and so are many ML algos)
- We had 16bit fixed point computations
 - DSPs (e.g. Motorola 56k)
 - Dedicated Neurocomputers (Synapse1)
 - Neurocoprocessors (Synapse3•PC) with hardware-independent programming framework – also auto scaling (Systolic Array)

FROM A MACHINE LEARNING POV FP16 SHOULD BECOME A STANDARD TYPE

FLASHBACK FROM THE PAST: SYNUSE•BASE DATATYPES WERE COMPLETELY HARDWARE AGNOSTIC

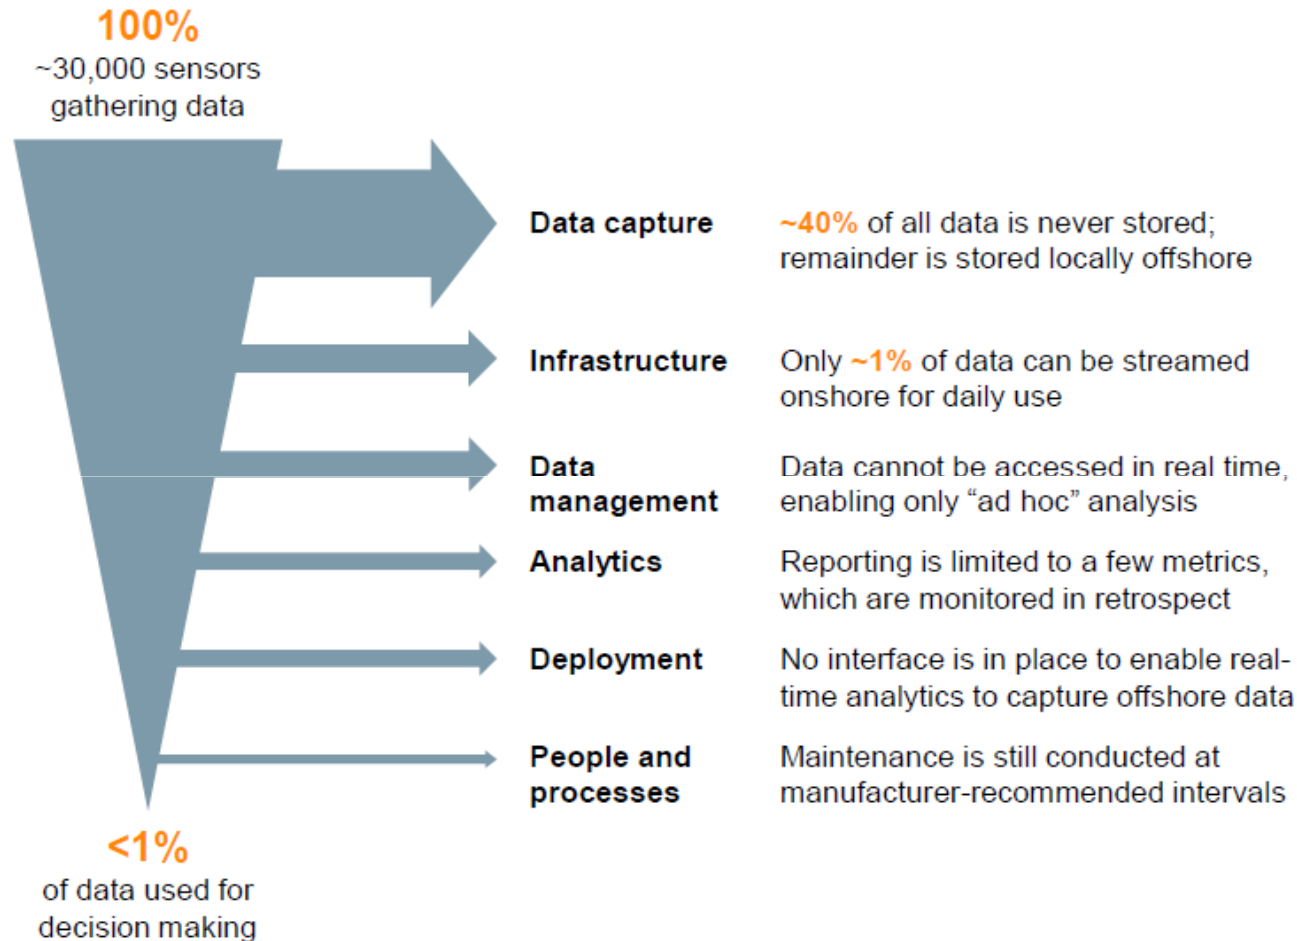


FP16 WHISHLIST: FULL HARDWARE INDEPENDENT SUPPORT

- Type support in common languages
 - `short float` float, double, long double
- Compute support in standard cpu
 - X86 via AVX(2) intrinsics
 - ARM via NEON (>8.2)
 - ...but not out of the (C++)-box
- Opaque transformation into GPU instructions
 - OpenCL (only for storage via extension)
 - CUDA?

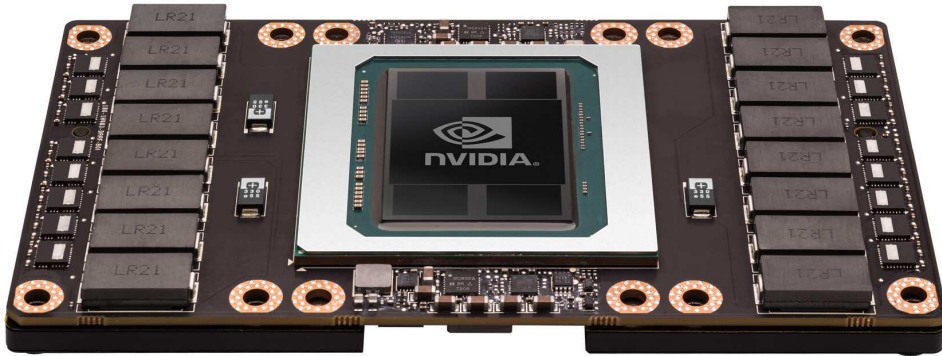
FOMITCHEV, ET. AL: „ADDING FUNDAMENTAL TYPE FOR SHORT FLOAT”

ONLY 1% OF DATA ON OFFSHORE LOCATIONS CAN BE USED REMOTELY



SOURCE: McKinsey Global Institute analysis

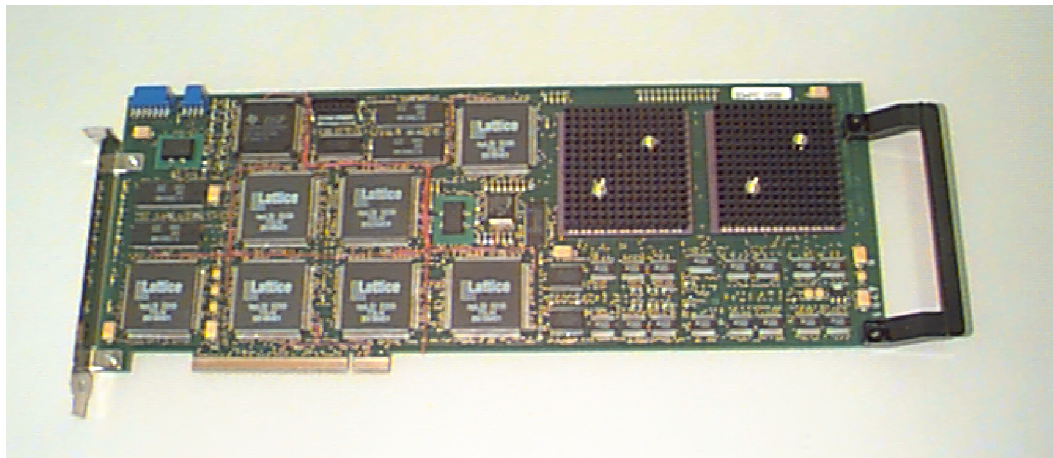
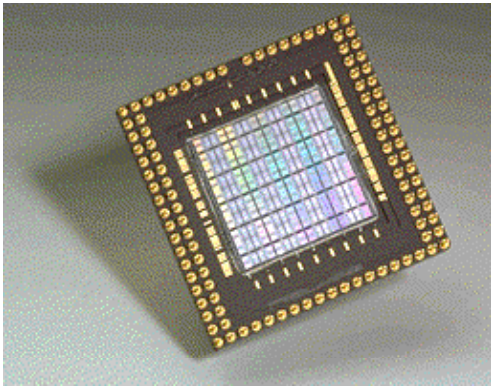
EDGE COMPUTING IS COMMODITY AND NECESSITY



LOCAL ENTITIES GENERATE TOO MUCH DATA TO BE TRANSPORTED TO CENTRAL LOCATIONS

- An autonomous car is expected to produce 1 TiB/h
- Full data set required to provide service
- Only partial data set needed for central processing/central services
- Edge nodes provide local service plus filtering for transport to backend (20TOPS DL@20W)

GPU IS HERE TO STAY – ALL OTHERS WILL GO THE WAY OF THE DODO



THE PAST HAS SHOWN, THAT HIGHLY SPECIALIZED CARDS DO NOT GAIN ENOUGH MARKET SHARE TO SURVIVE

- Wo still knows
 - CNAPS
 - Synapse (OK, now you know a little)
 - Intel Ni 1000
 - Intel ETANN (80170NX)
 - ...
- But we still have
 - X86
 - POWER
 - ARM
 - ... all now with multicore

THANK YOU!