Application Acceleration Engine Using GPGPU

Accelerative Technology Lab

Innovation for Life™

About Mimos

- **Malaysia's National R&D Center**
- **10 core research areas:**
  - Advanced Analysis & Modelling
  - **Advanced Computing**
    - ***Accelerative Technology Lab***
  - Information Security
  - Intelligent Informatics
  - Knowledge Technology
  - Microenergy
  - Microelectronics
  - Nanoelectronics
  - Psychometrics
  - Wireless Communications

- **Advanced Computing**
  - Spearheads R&D activities in acceleration on large-scale computing, chiefly Cloud Computing; from SaaS and IaaS to Services Delivery Platform.

Open Platform

Traditional Multicores
(Main Processor)

General Purpose GPU
512-2880 GPU Cores
(Co-Processor)

Many Integrated
CPU Core (60
Cores)
(Co-Processor)

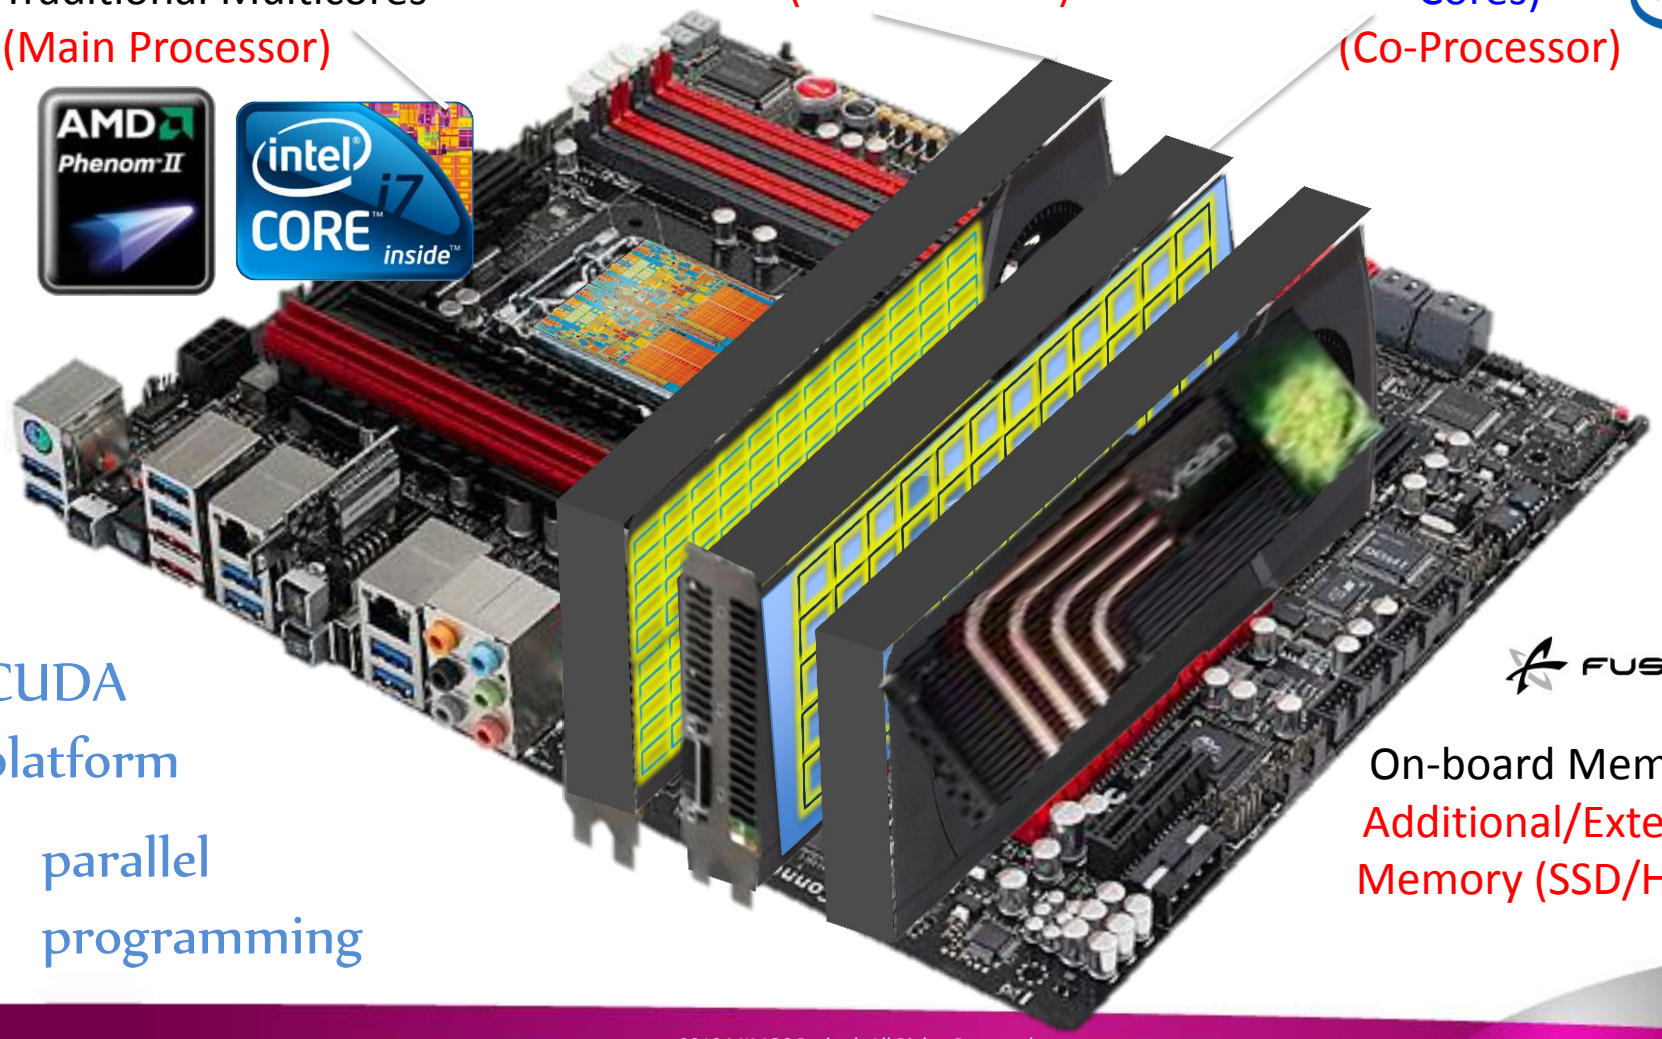

CUDA
platform

parallel
programming

On-board Memory
Additional/External
Memory (SSD/HSM)

GPU/Multi Core Driven Applications for MIMOS

Large DataSets (SOCSO)

Streamed Data (ISP & AVMedic)

En/Decryption (scrambling+) + SSL Accelerator

MiMORPHE+

Data Processing

Accelerated Processing

PRDRM (HRMIS) + Patient Data Analytics

(BigData)

Business/ Enterprise Data

(BI)

Fraud Mgnt System (SOCSO)

MiAccLib 2.0 Release

AccLib

Legend

- 1st Version Done
- In Plan/ Progress

Selected GPGPU/MiAccLib Projects

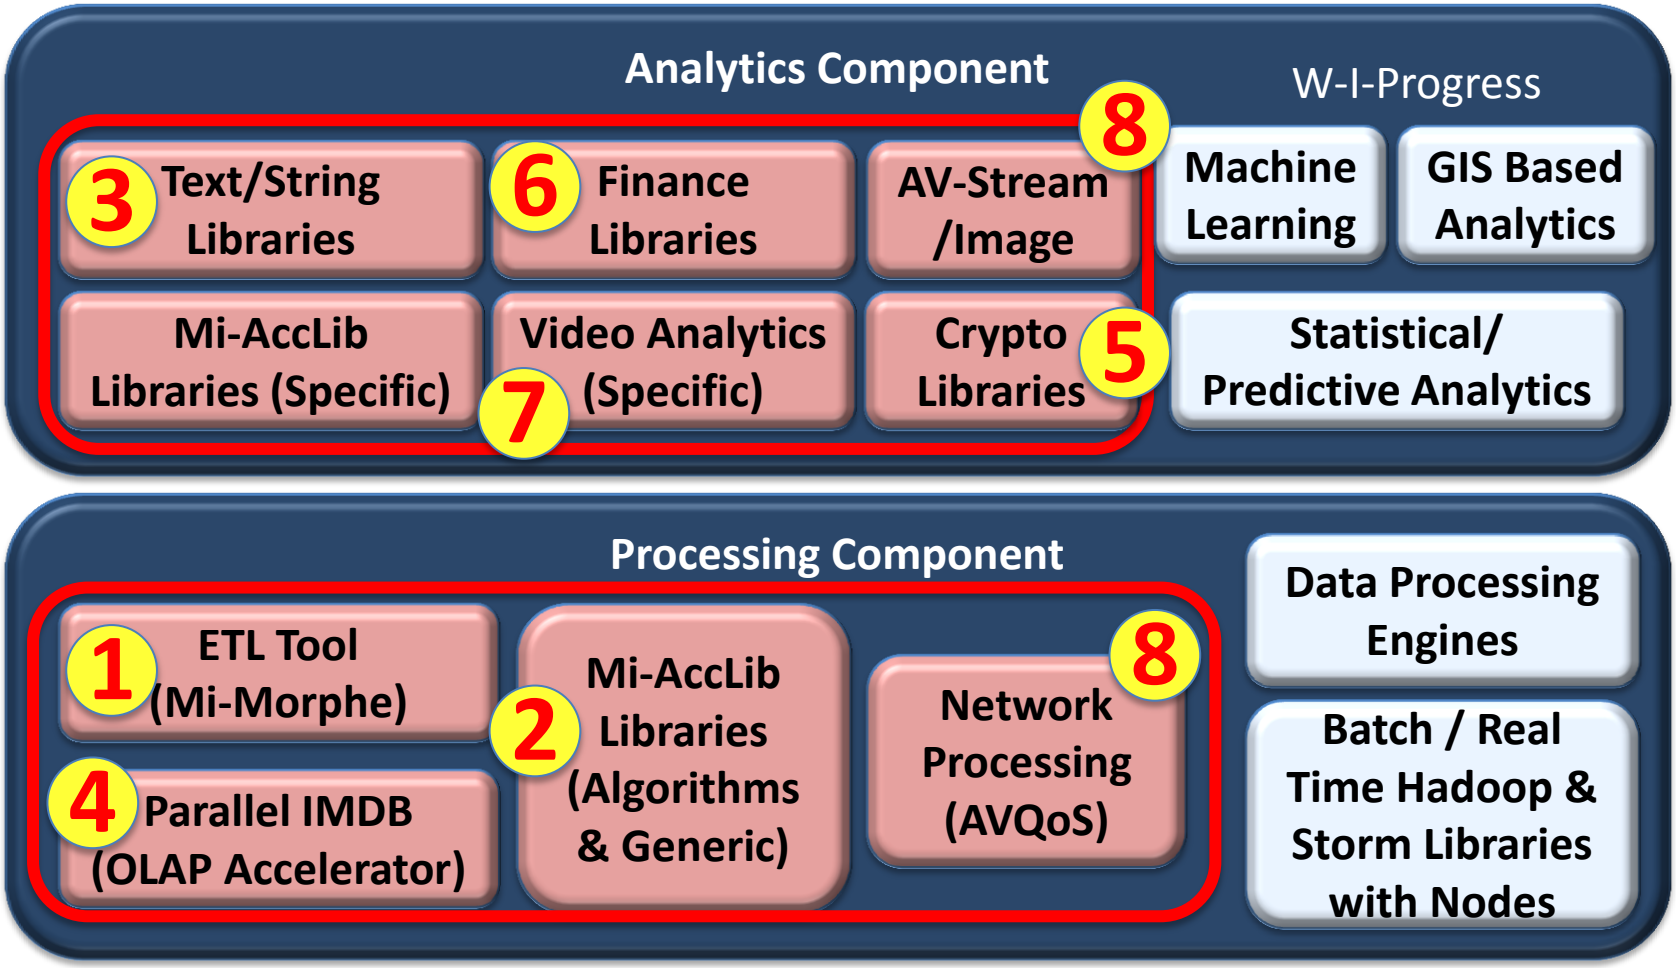

MIMOS Accelerated Library

AccLib

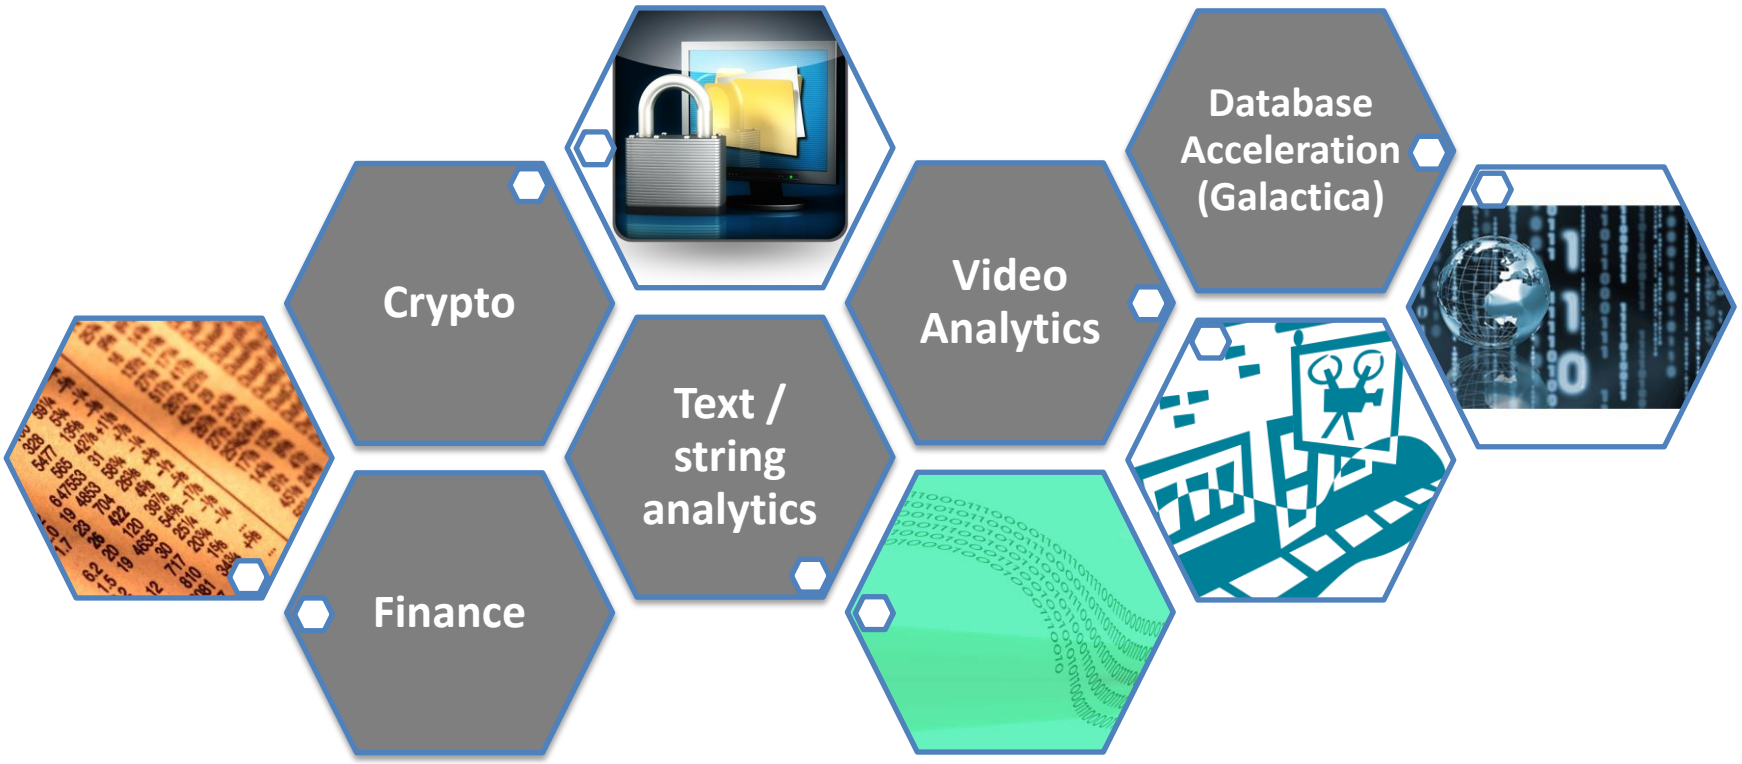

Database Acceleration

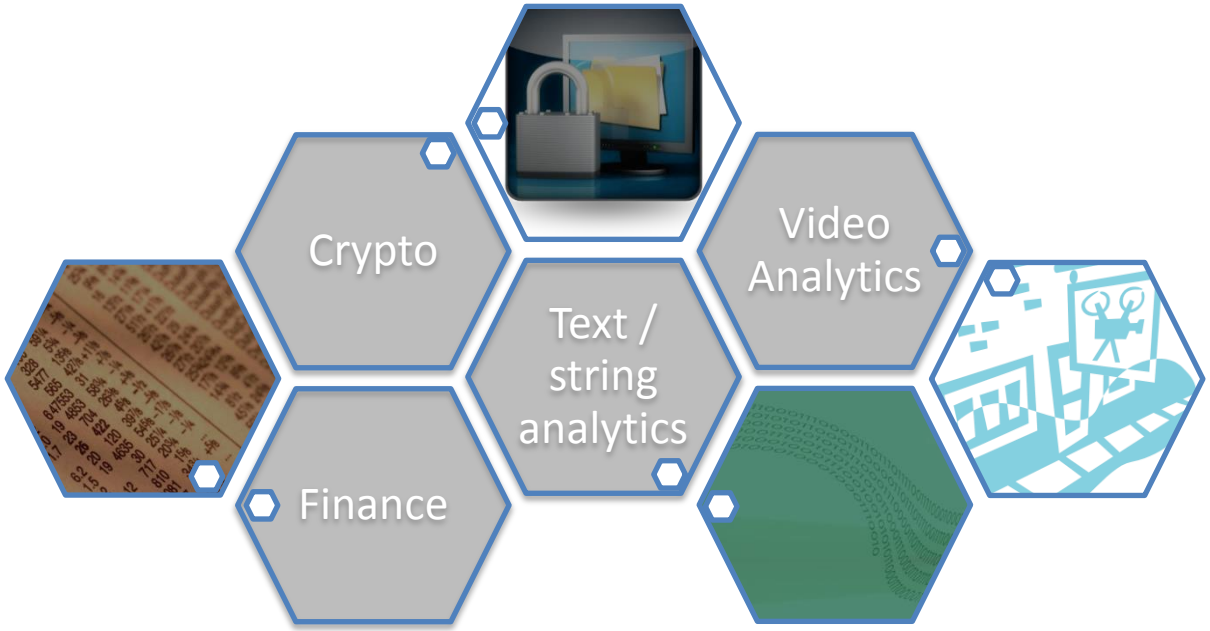

Research for Queries Acceleration

Vision

Capitalize GPU parallel architecture for large data processing via queries Parallelizer

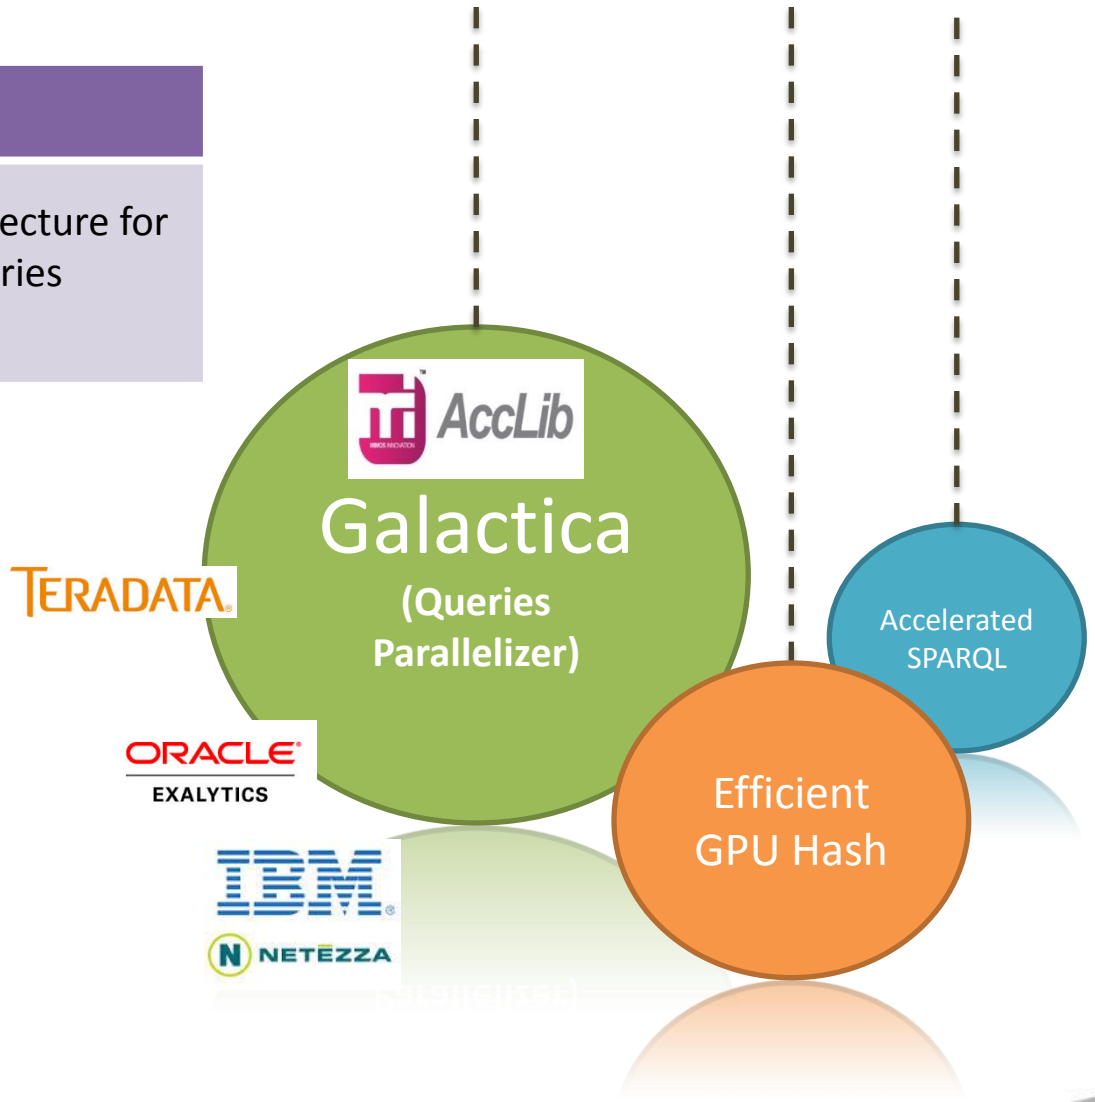

GPU-based Database Acceleration (Galactica)

GPU Parallelization

Enabling Vector Processing
Parallel Compression in GPU
Data Crumble
Stream Processing
SIMD and MIMD

GPU Query Processing Engine

Selection, Join, Aggregation
Sorting, Searching
Meta Analyzer
Queries Optimizer

Support SQL

Functions:
COUNT, SUM, AVG,
MIN, MAX, LIMIT,
DISTINCT

Operators:
!, =, <>, <, >, <=, >=, <=>,
<<, >>, |, &, +, -, *, /, %,
^, MOD, AND, EQUAL, OR, XOR, NOT

Data Types

Integer
Float
String
Date

Result of *MiAccSQL* vs *PostgreSQL*

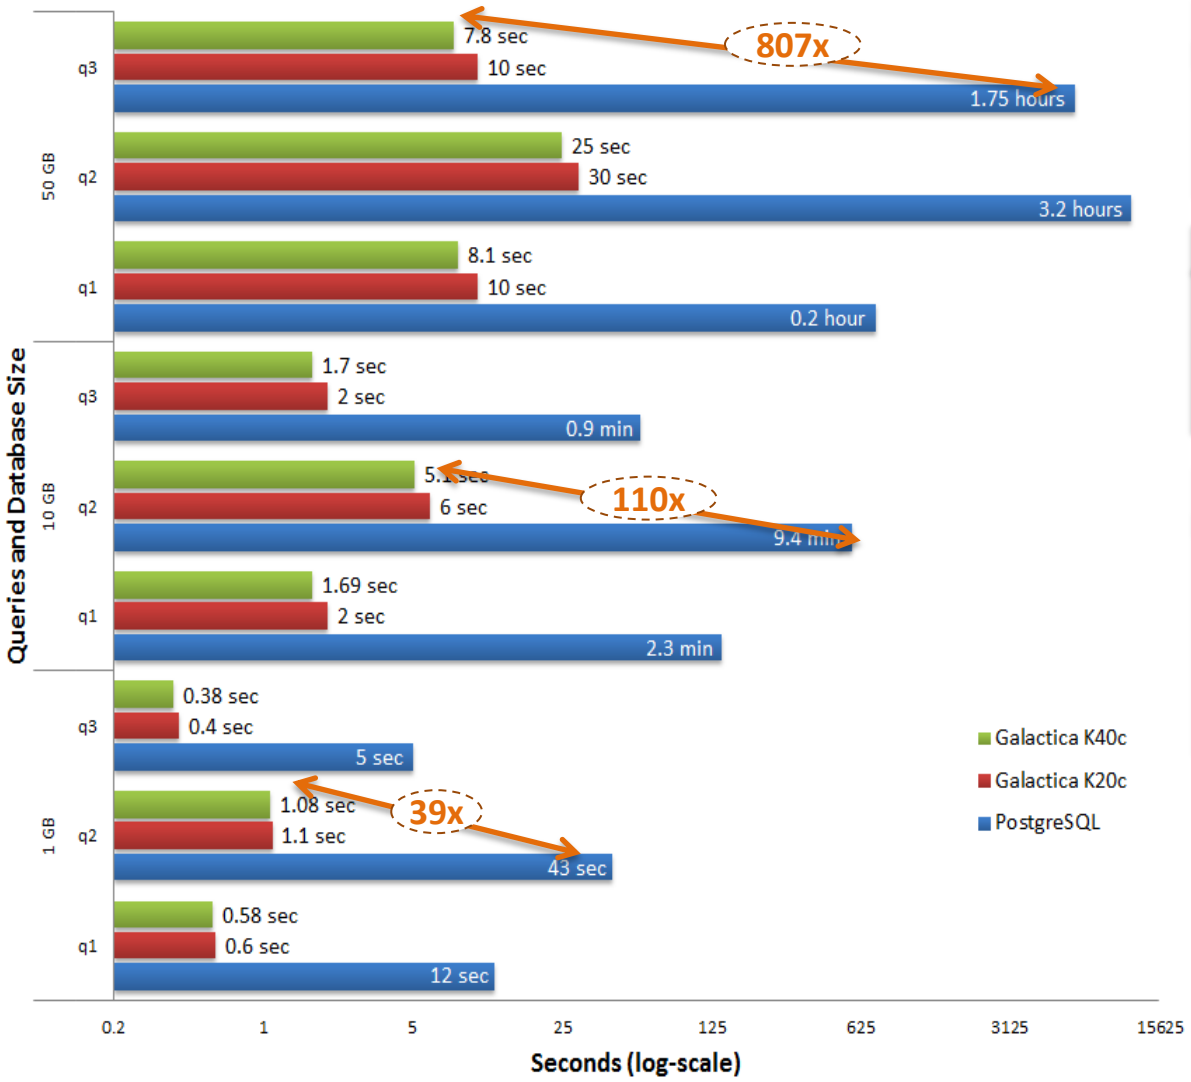

***Data is taken from TPC-H benchmarking**

SQL Query 1

Compute amount of business that was billed, shipped and returned

SQL Query 2

Compute the total revenue, quantity and orders from the "Building" customer

SQL Query 3

Compute the revenue, total and average amount of quantity along with the average price from transactions

** Setup Config:

CPU	Intel(R) Xeon(R) CPU X5680 @ 3.33GHz (2 processors)
RAM	22 GB
GPU	NVIDIA Tesla K40c / K20c
CUDA	5.5
Storage	WD HDD 1TB
Database	PostgreSQL 9.3
OS	Windows 7 (64 bits)

Result of Galactica *vs* Hadoop

32GB Data

- Best performing on sum
- Cost saving HPC
- Failed queries operation because Galactica does not support the feature yet

Hadoop	GPU
7 Virtual Machines with one master node (8 cores)	DELL Precision T5500 workstation
6 worker nodes (4 cores each) running on a few of	NVIDIA Tesla K20c
HP DL380p G8 servers installed with Apache Hadoop,	on Intel Xeon E5630@2.53GHz processor
Cloudera's Hadoop and Impala.	12GB RAM
Postgres on another same model of HP server with 8GB RAM with 4 cores and another high end HP machine with 96 GB RAM and 48 cores.	1 TB Sata Hard drive (7200rpm).
	Windows Server 2008 R2 Enterprise SP1 64-bit

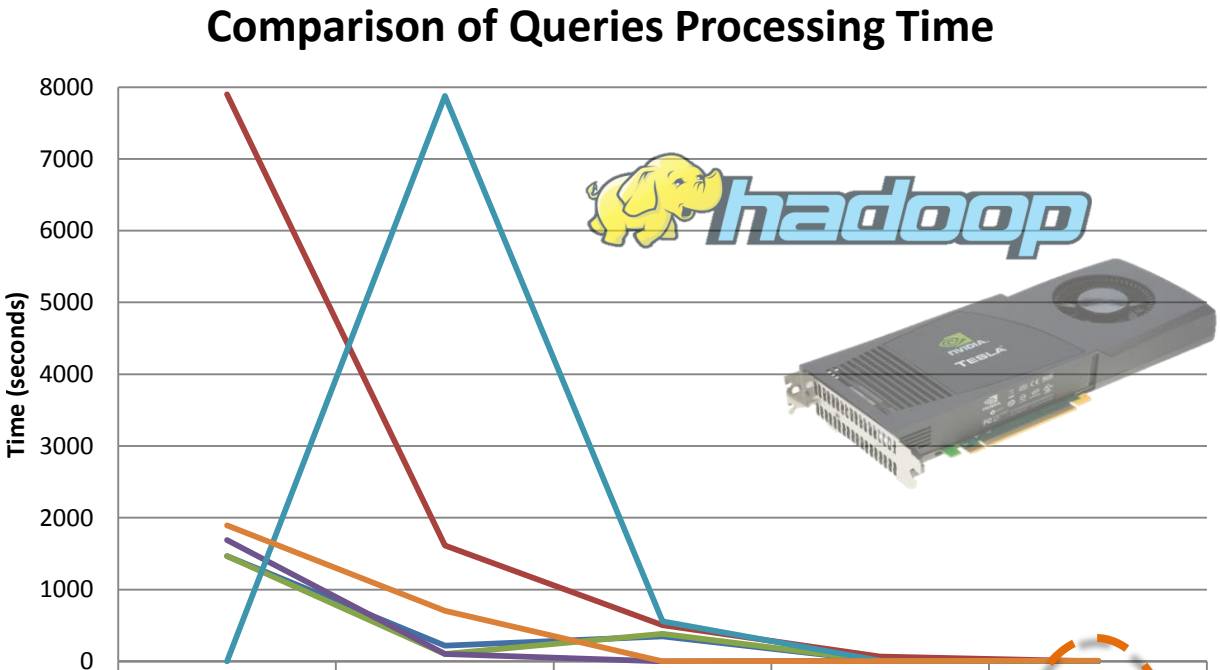

	SQL (8GB/4Core)	SQL (96GB/48Core)	Hadoop-Hive	Impala-Hive	Galactica
Q1	1466.7	218.7	347.6	3.7	0.3
Q2	7901	1612	505	64.2	N/A
Q3	1464.7	103.6	383.5	3.5	3
Q4	1688.7	102.7	N/A	2.9	1.6
Q5	Failed	7878	557.3	10.1	N/A
Q6	1893	704	N/A	3.7	6.3

Text / String Analysis

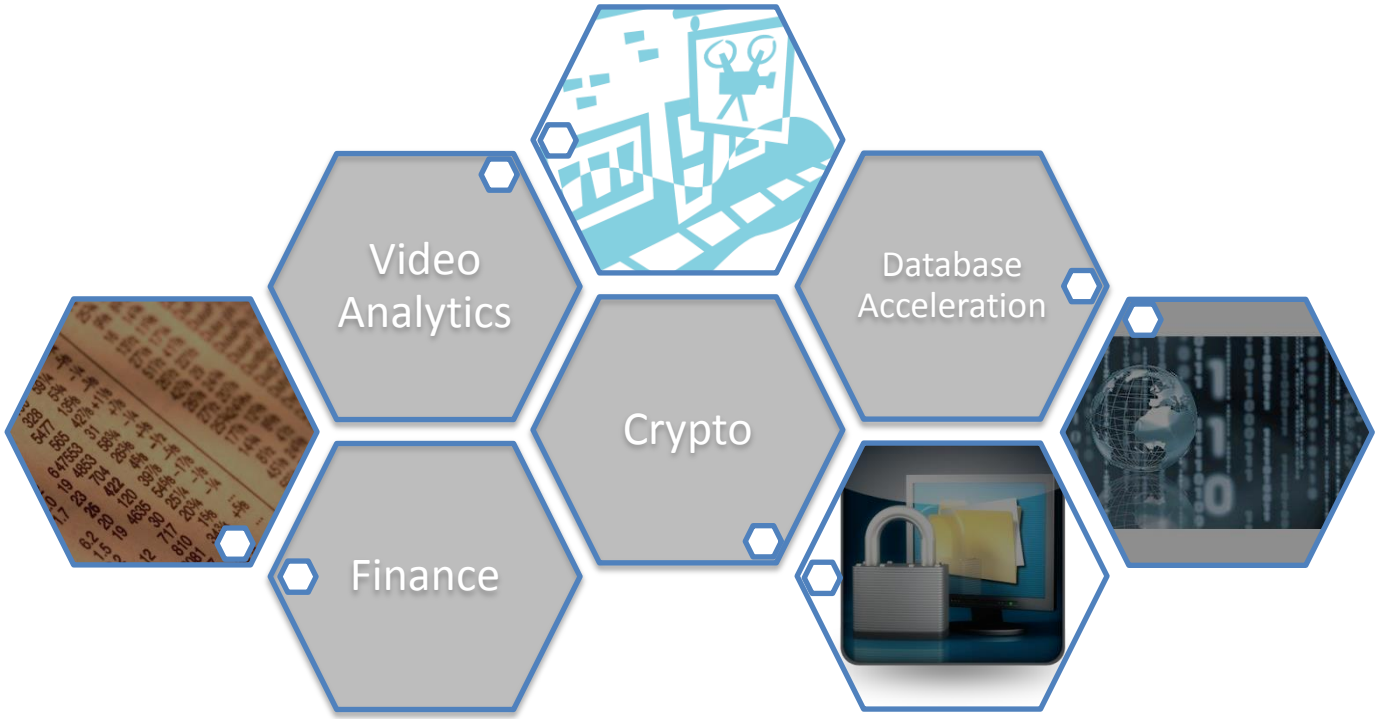

High-Speed Name Search Performance

Mi-AccLib

Accelerated & Parallelized Algorithms

10+ Million Records of transaction data

USE CASE

Perkeso Data X JPN Data

350 Trillion
X 7711 Rules combinations/rule

2.7 Quintillion Operations

Exact Match	Mohamad	Mohamad
Edit Distance	Mohamad	Muhammad
Wild Card	Moh	• Mohamad • Lee Ang Moh • El-Mohan

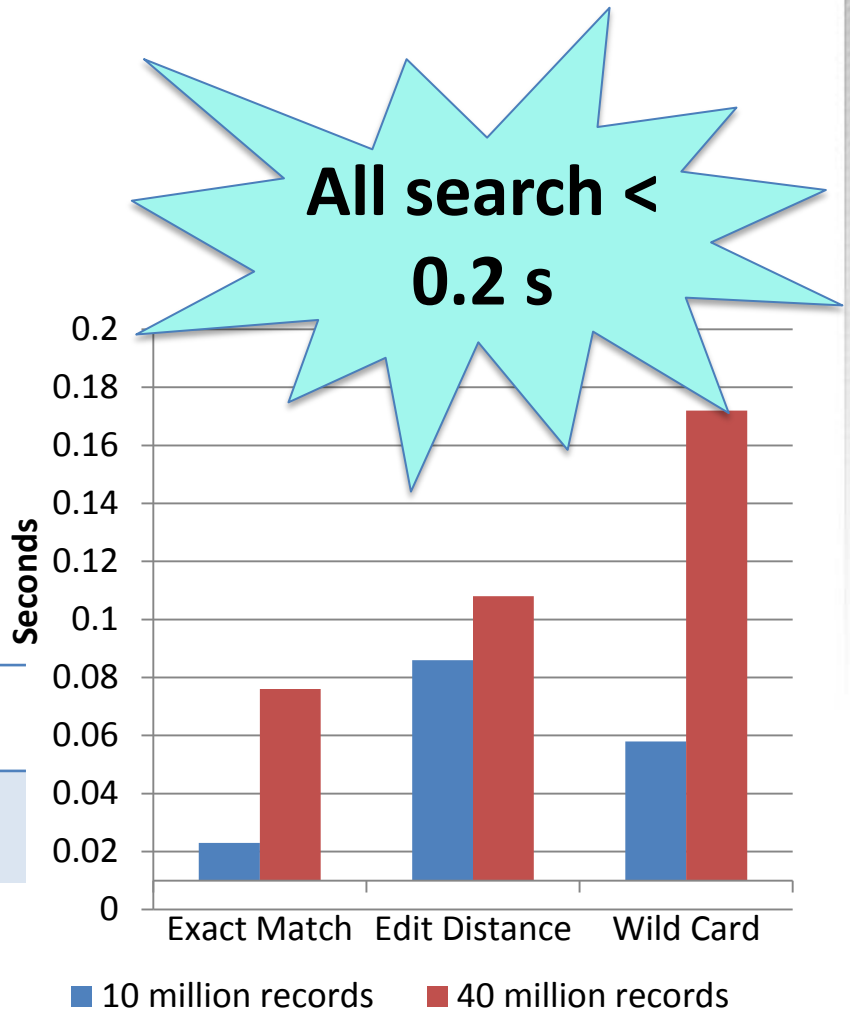

Use Case: SOCSO

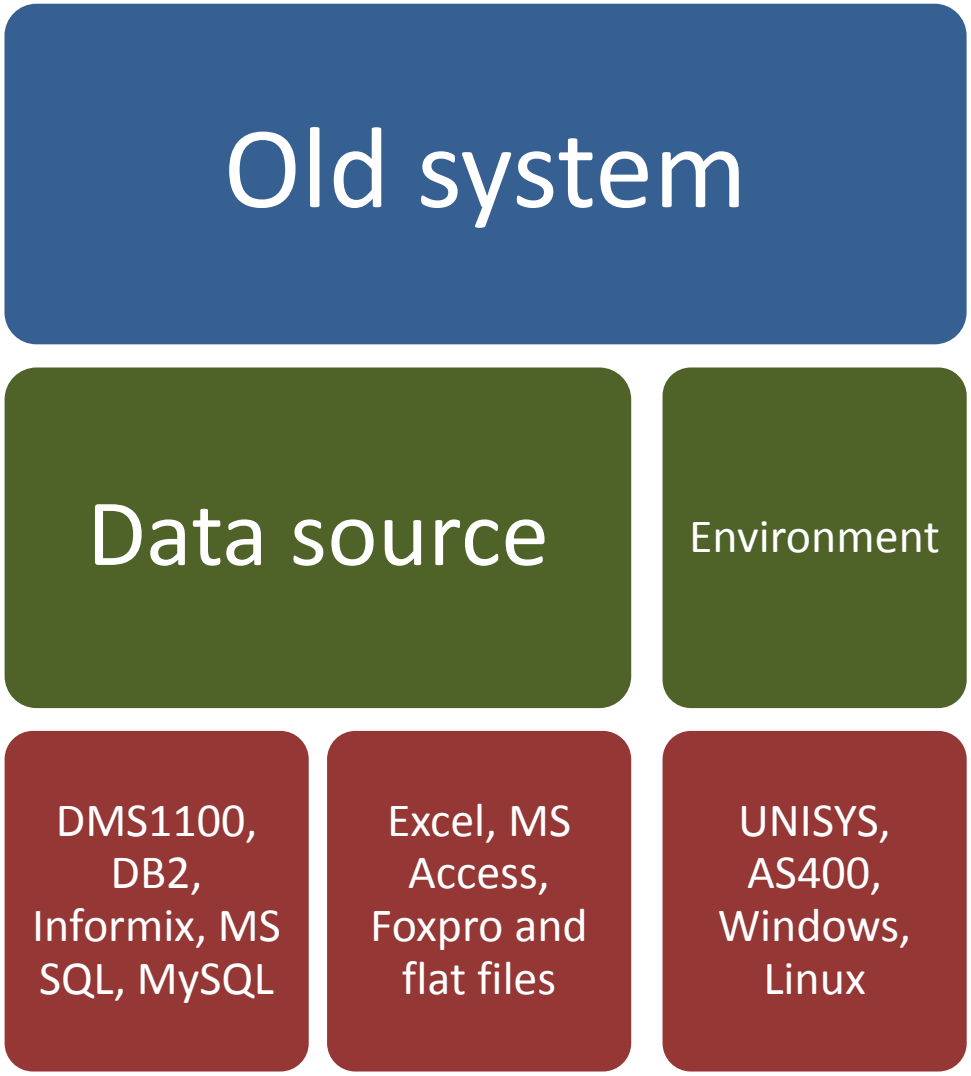

**** 7711 cleansing rules**

Big Data:

- **319 source data**
- Involves ~**1 billions records**, e.g.:
  - 15 millions employee with 150 millions of monthly contribution
  - 880,000 employer with 65 millions of monthly contributions
  - Match against reference JPN data with 15 million records

Mi-Morphe Data Flow/Process

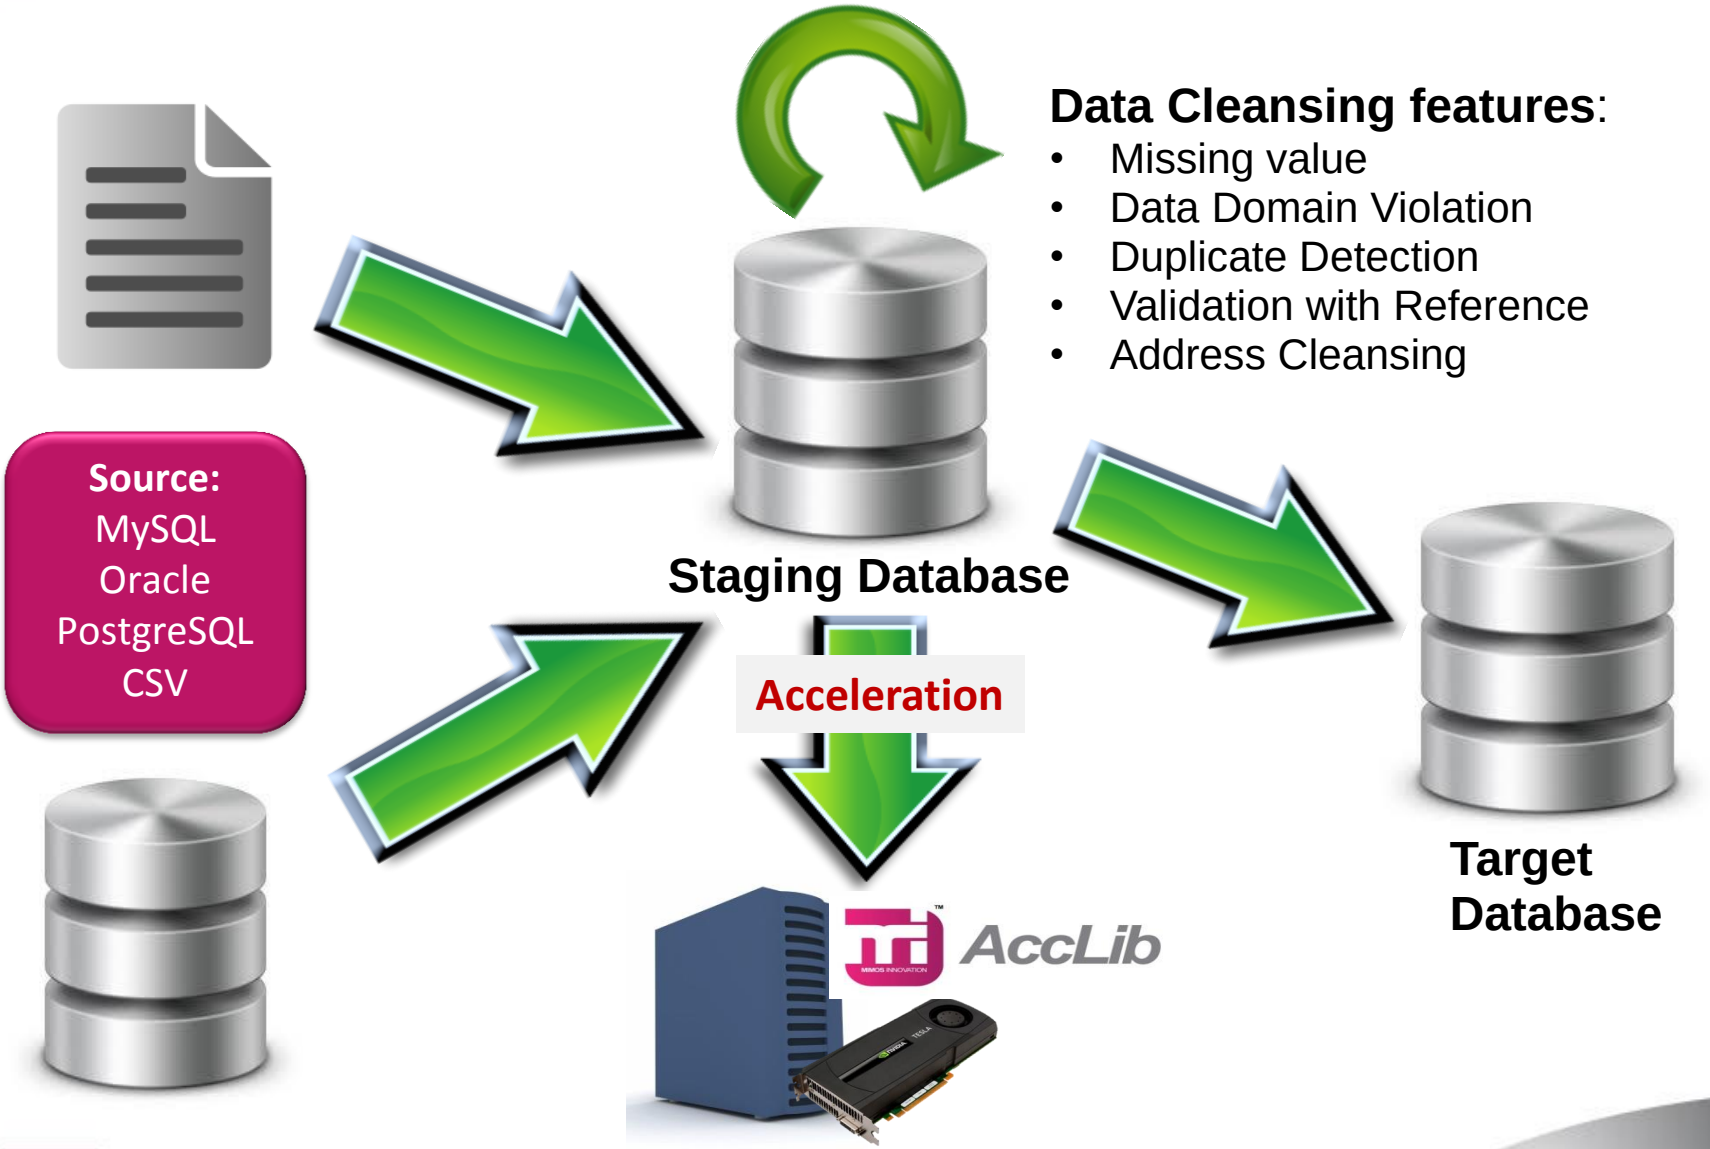

Data Cleansing features:

- Missing value
- Data Domain Violation
- Duplicate Detection
- Validation with Reference
- Address Cleansing

Mi-Morphe Data Cleansing Algorithm/Feature

No	Feature	Remark
1	Name Comparison	Detect similarities of name based on different Names using **MiAccLib**
2	Duplication detection	Detect duplication based on **Edit Distance**, **Soundex**, **Numeric distance, Date distance**, **Q-grams** and **Levenshtein ratio** algorithm using **MiAccLib**
3	Record Linkage	Verify record is reference table based on **Edit Distance**, **Soundex**, **Numeric distance, Date distance, Q-grams** and **Levenshtein ratio** algorithm using **MiAccLib**
4	Address harmonization	Detect address abbreviation base on address ontology using **MiAccLib**
5	Address field chunking	Auto chunk address based on country, state, post code, district
6	Auto/Batch correction	Allow auto batch correction for bunch of records using **MiAccLib**
7	Missing value detection	Detect mandatory field violation
8	Data domain detection	Detect orphan record using **MiAccLib**
9	System Assisted Manual Rectification	User friendly UI for data assessment and rectification

Condition Detection for Employee - JPN Validation

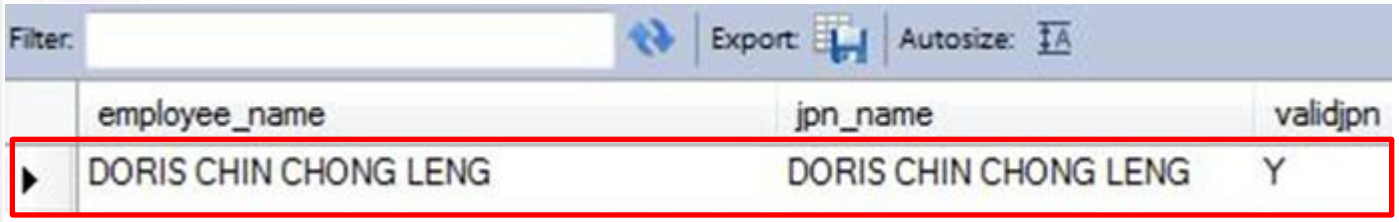

Algorithm Used

Exact Matched

NOORLAILY BT SAMSUDIN	NOORLAILY BINTI SAMSUDIN	Y	**Subtracted Matched**
SHU QIN TAN	TAN SHU QIN	Y	**Words Check**
MAU BOO BT DASTAGIR	MAU BOO BEE BINTI DASTAGIR	Y	**Levenshtein Ratio**
ABDILLAH B HJ RASOL	ABDILLAH @ ABDULLAH BIN ABDUL RASOL	Y	**WCSnd & WCLR & FLName**
SUZIE ARIANTI BT YASIN	LEE KWAI MOY	Y	**System Assisted Manual Verification**

350,000,000,000,000 (350 Trillion) combinations per complex detection rule

Address Harmonization Complexity

Detect, correct, transform and harmonize address value to pre-defined format and abbreviation.

Example

Selangor 61 ways written!	slngr;slgr;s'ngor;s'lgr;slgr.;sel.;sel;s'glr;s.d.e;s'gor;sgor;selgr;sngr;sgr;selasngor d.e.;sngor;selasngor de.;drlsngor d e;slngor;selangr;selangro;selngr;slangor;slangor d e;slangor d. e.;slangor d/ehsan;slangor.;selngor;selngor d e;selngor.;selngor d e;selngor darul ehsan;sgor d.e.;sgor de;selangor darul ehsan;selg; selgor; sel. d. e.; slgor;selangor d e; selangro; selangro darul ehsan; selangro d e.; selangr d/ehsan; selangr d/e; selangror d e; selangror darul ehsan; selangor de; selangor, darul ehsan; selangao; elangor d e; s'ngor d.e.; selangor d.ehsan.; sel d e; selangor d.e,; selanngor.; selangor de.; sel. d.e.; seelangor; selangor darul ihsan; selangor darul ehsan.;
Kuala Lumpur 41 ways written!	kuala lumpor;k lumpur;k. lumpur;k.lumpur;kl;k.l;kuala-lumpur;kuala lumpure;kuala lumpuyr;kuala luumpur;wilayah persekutuan kuala lumpur; k l; kuala lummpur; k.l.; k lumpur.; k.lumpur.; kuala lumpur0; kuala lunmpur; k. l.; kuala umpur; kualalumpur.; w.p k/lumpur.; kualalumpur; k lumpures; 50050kl; ku la lumpur; kualu lumpur; kuala lumpur wilayah persekutuan; w . p kuala lumpur; w . p . kuala lumpur; w. persekutuan kuala lumpur; w. persekutuan kuala lumpur.; w.p kuala lumpur; w.persekutuan k l; w.persekutuan kuala lumpur; wilayah persekutuan k.l; wilayah persekutuan k.lumpur; wilayah persekutuan, k.l; wilayah persekutuan,k.l; wilayah persekutuan. k.l; wp kuala lumpur;

AccLib Data Cleansing Performance

-MORPHE

RAW DB in CSV Format

NVIDIA CUDA

Results in CSV Format

Data Detection:

1. Exact match
2. Edit distance

~15 Millions

vs **~15 Millions**

search

Duplicate Records Performance Comparison Using Edit Distance Algorithm

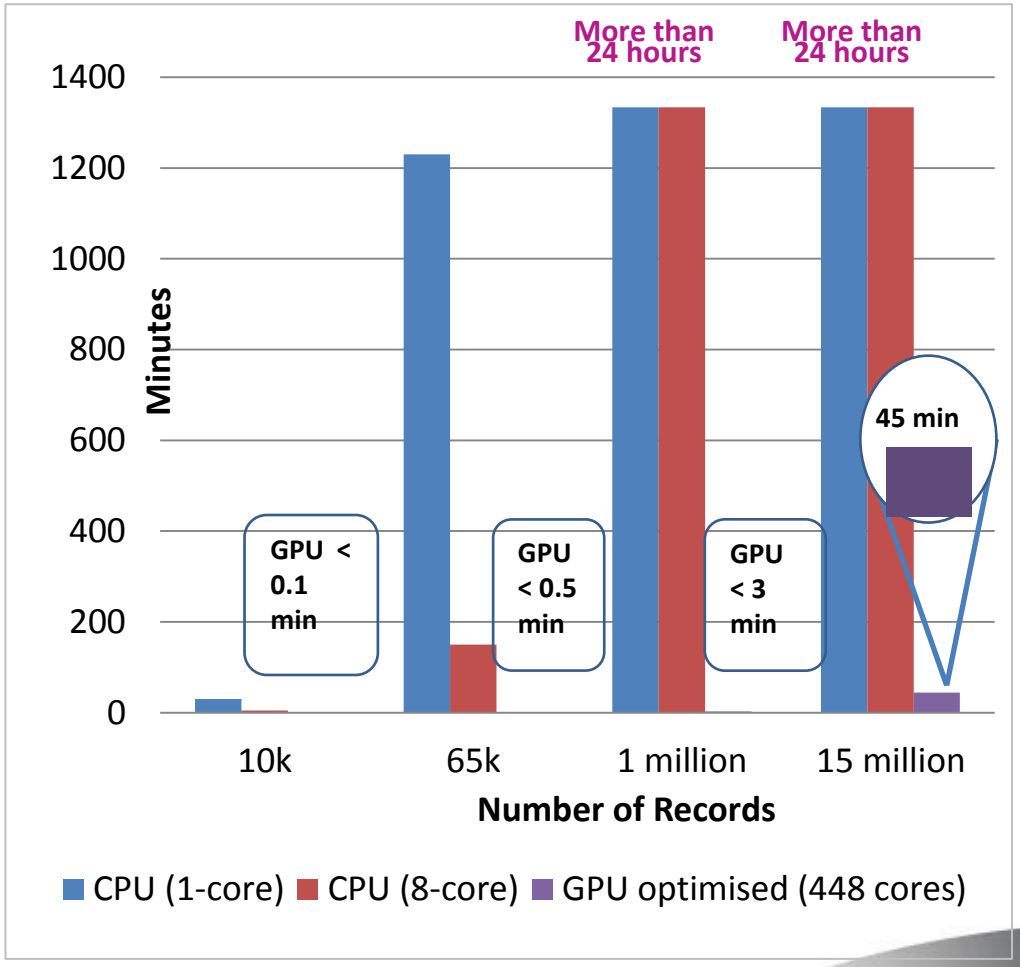

Twitter Analysis With MiAccLib

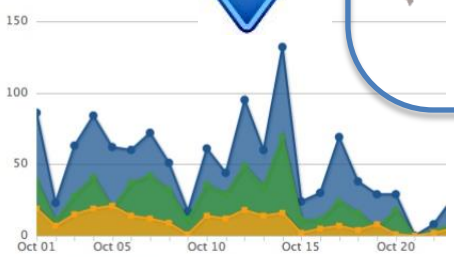

AccLib

NVIDIA CUDA

Keyword Occurrence Search (Wild Card)

New

Wild Card Search comparison between CPU and GPU

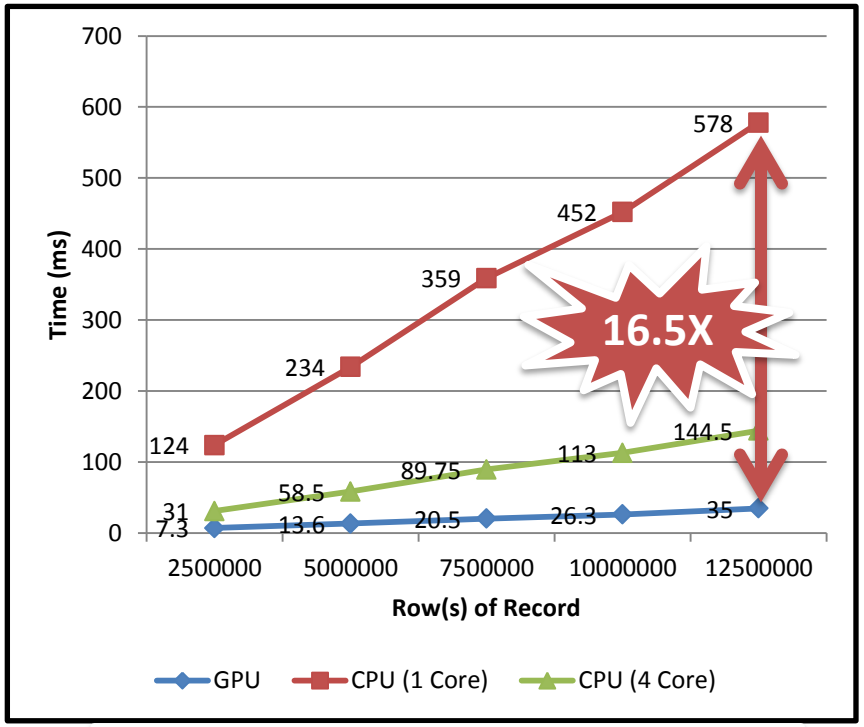

16.5x Speed Gain!!!

Target Customer: (confidential)

Text/String Matching Algorithm on Two GPU Cards

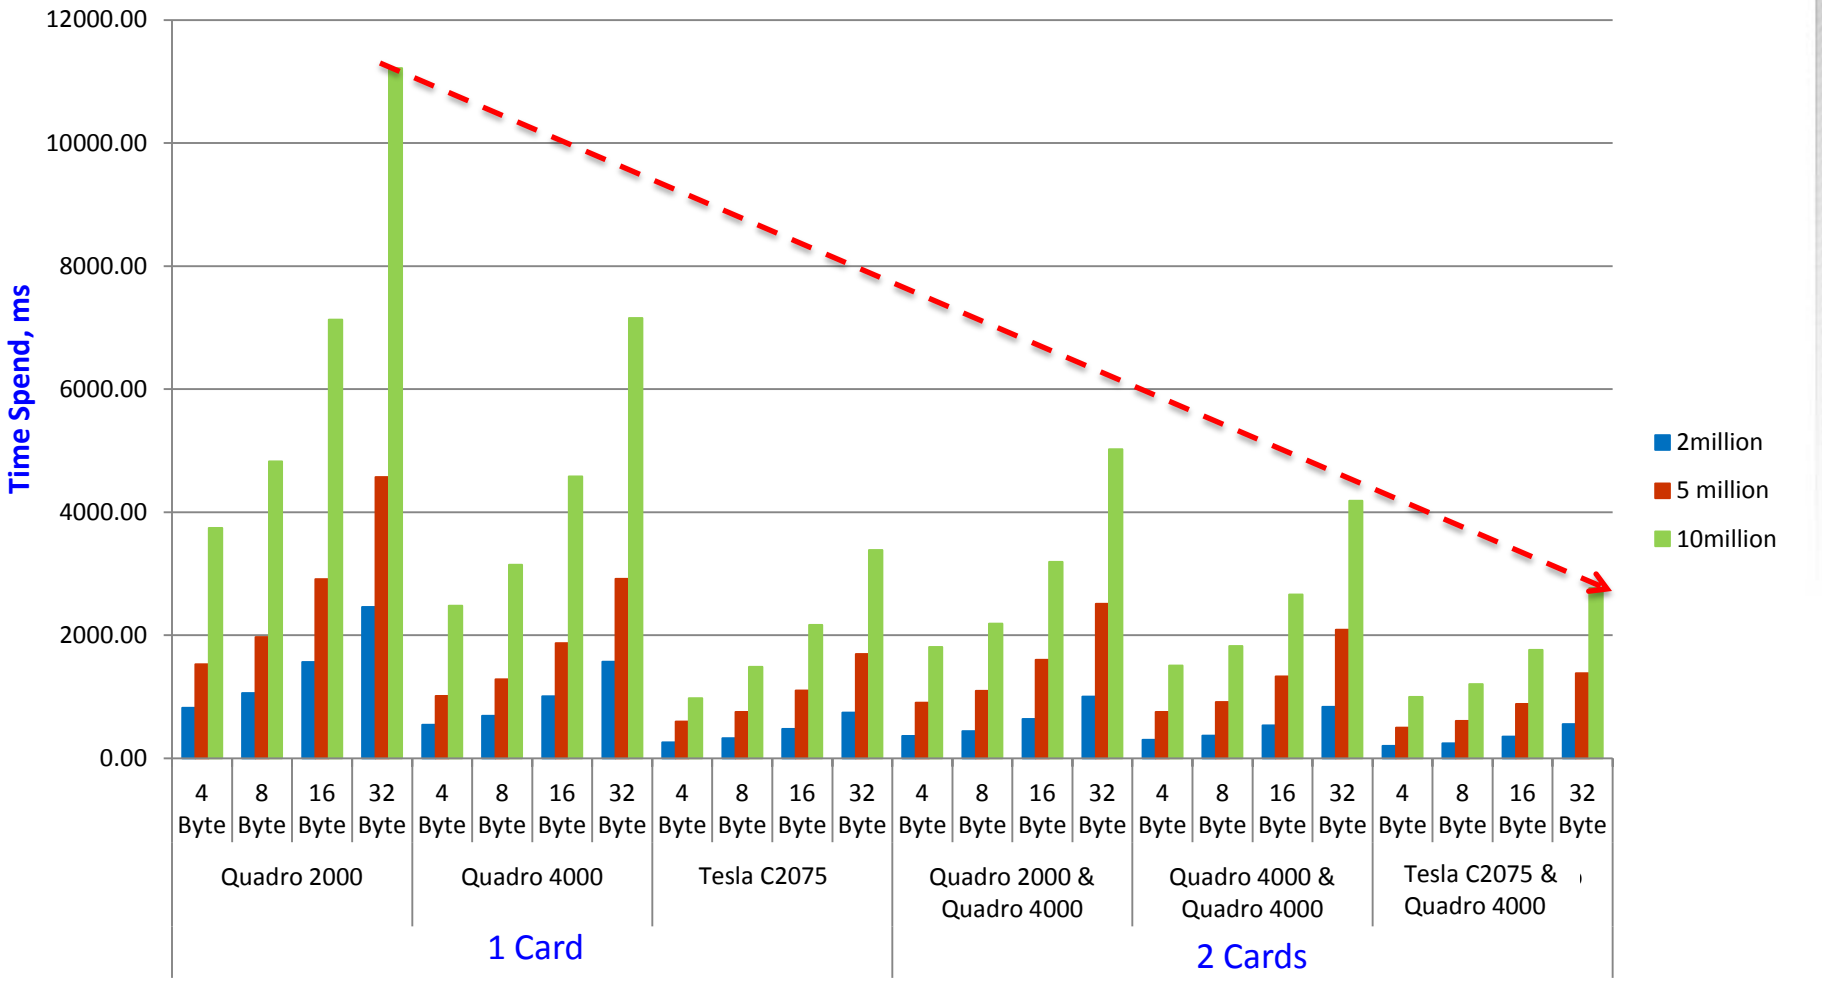

HDD vs SSD for Various Data Sizes in 'Preprocessing' Stage

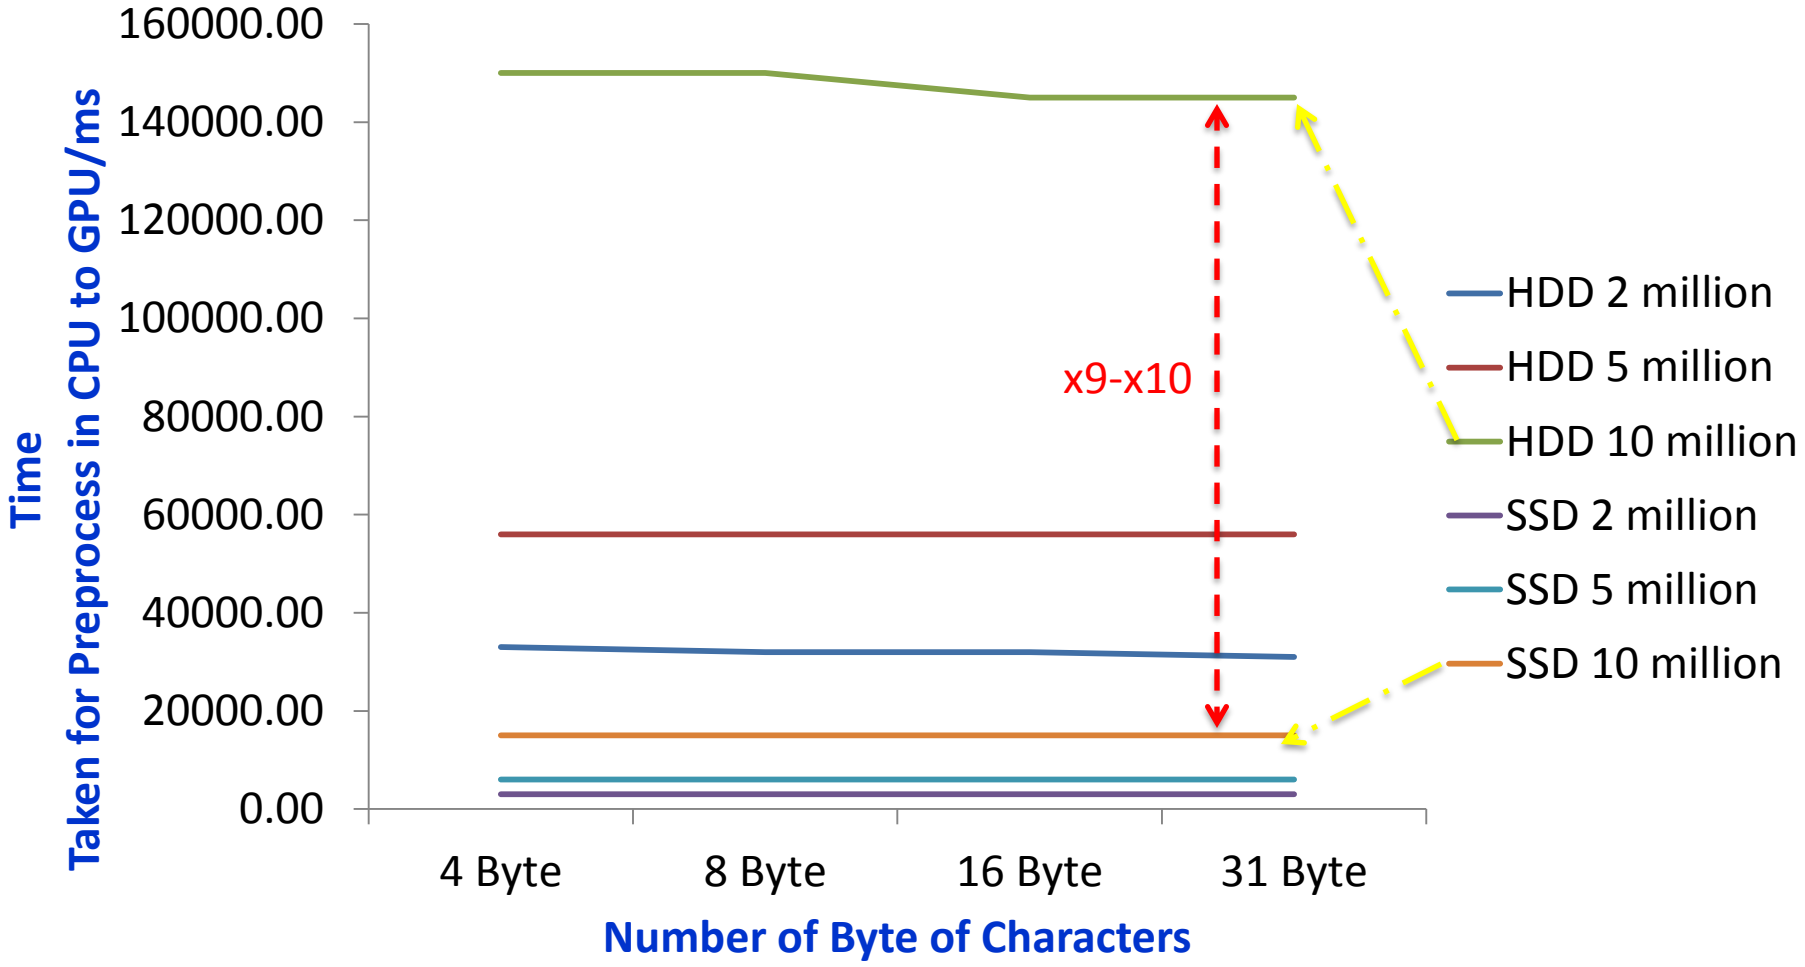

Complexity Edit Distance/Levenshtein Distance

Mathematically, the Levenshtein distance between two strings a, b is given by $\operatorname{lev}_{a,b}(|a|,|b|)$ where

$$\operatorname{lev}_{a,b}(i,j) = \begin{cases} \max(i,j) & , \min(i,j) = 0 \\ \min \begin{cases} \operatorname{lev}_{a,b}(i-1,j)+1 \\ \operatorname{lev}_{a,b}(i,j-1)+1 \\ \operatorname{lev}_{a,b}(i-1,j-1)+[a_i \neq b_j] \end{cases} & , \text{ else} \end{cases}$$

Note that the first element in the minimum corresponds to deletion(from a to b), the second to insertion and the third to match or mismatch, depending on whether the respective symbols are the same.

Mathematical Formula

Example

For example, the Levenshtein distance between "kitten" and "sitting" is 3, since the following three edits change one into the other, and there is no way to do it with fewer than three edits:

kitten → **s**itten (substitution of "s" for "k")
sitt**e**n → sitt**i**n (substitution of "i" for "e")
sittin → sittin**g** (insertion of "g" at the end).

Socso Scenario

Ahmad -> Ahmat
Samy -> Sami
Ah Moi -> Ah Moy

Example calculation on Text Data

		B	O	B	S
	0	1	2	3	4
B	1	0	1	2	3
O	2	1	0	1	2
B	3	2	1	0	1
B	4	3	2	1	1
Y	5	4	3	1	2

Finance

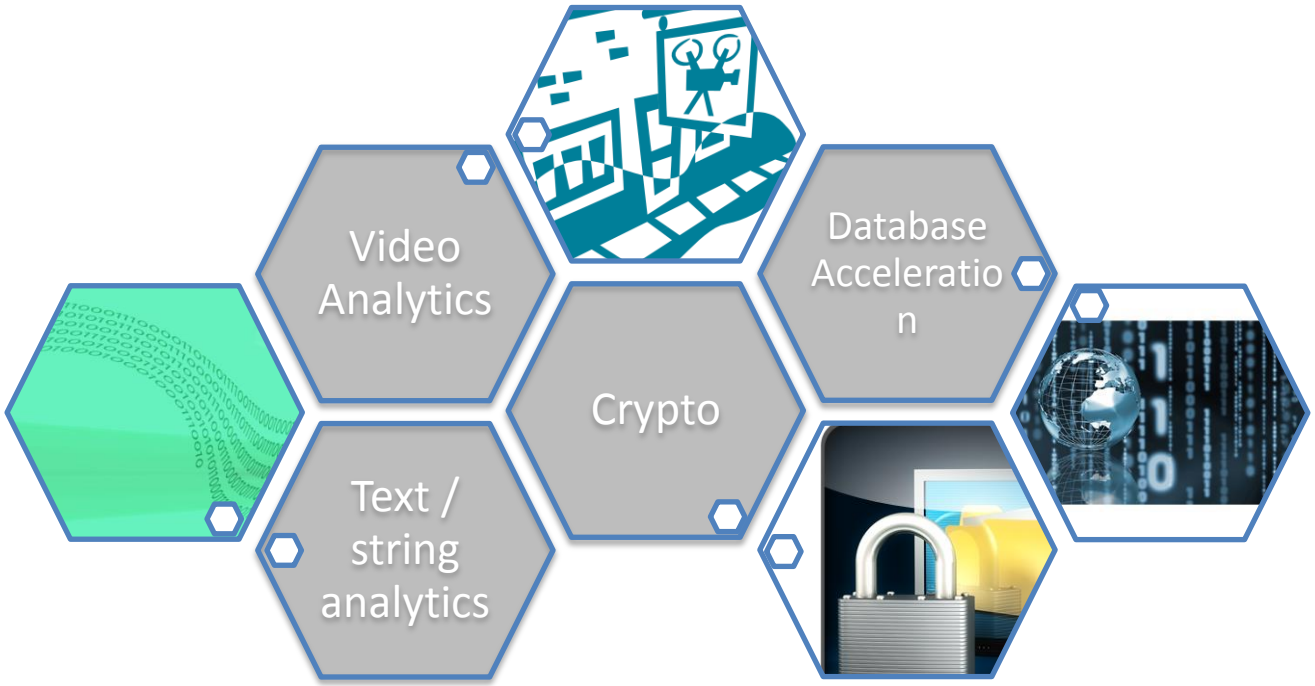
Video Analytics
Database Acceleration
Crypto
Text / string analytics

Financial Data Analytics

TES·S INTERNATIONAL

Generalized & Parallelized Algorithms

CORAL iSEM™ Intelligent sensor for money laundering

- Anomaly Analysis
- Abnormally Transaction Prediction Algorithms

Millions of Transaction Screening

Intelligent Money Processing for Anti Laundering

10+ Million Records

3.5 Millions Financial Terrorist to be checked

Moving from Hours to Minutes/Seconds

Legend

- 1st Version Done
- In Plan/ Progress

Financial Analysis With MiAccLib

Stock Pair Trading

High Correlation Stock Pair

Pair Correlation = $$\frac{\sum(A_i-\overline{A})(B_i-\overline{B})}{\sqrt{\sum(A_i-\overline{A})^2\sum(B_i-\overline{B})^2}}$$

- Provides the user a historical view on the correlation between different pairs of shares.
- Shares with high correlation historically move in the same direction.

Multiple Portfolios with Mutliple Stocks

Historical Data for Selected Stocks (Daily)

Trading Parameters (e.g):

- Ratio versus Spread
- Moving Average
- Standard DeviatioΩ
- Maximum day in trade
- Stop Loss
- Critical Entry and Exit
- Start and End Date

Accelerated Calculation:

Price Ratio (**PR**)/Spread (**PS**)
PR/PS Moving Average
PR/PS Moving standard Deviation
Normal Deviation (**ND**)
Average Spread **Correlation**
Spread **Co-integration**

Compute Intense & Parallelizable Algorithms!

Buy/Sell/Hold Parallel Stocks in Different/Multiple Portfolios

Complexity of Computation for Pairwise Correlation

$$\text{Pair correlation} = \frac{\sum(A_i-\overline{A})(B_i-\overline{B})}{\sqrt{\sum(A_i-\overline{A})^2 \sum(B_i-\overline{B})^2}}$$

$$= \frac{(A_1-\overline{A})*(B_1-\overline{B})}{\sqrt{\sum^{2500}(A_1-\overline{A})^2 * \sum^{2500}(B_1-\overline{B})^2}} + \frac{(A_2-\overline{A})*(B_2-\overline{B})}{\sqrt{\sum^{2500}(A_2-\overline{A})^2 * \sum^{2500}(B_2-\overline{B})^2}}$$

$$+ \ldots\ldots\ldots\ldots + \frac{(A_{2500}-\overline{A})*(B_{2500}-\overline{B})}{\sqrt{\sum^{2500}(A_{2500}-\overline{A})^2 \sum^{2500}(B_{2500}-\overline{B})^2}}$$

Note:

2500 = 250 trading days * 10 years

319,600 = 800 stocks pair combination = $\frac{800*(800-1)}{2}$

For 1 pair:

Complexity = 2500^2

For 319,600 pairs:

Complexity = 319,600 * 2500^2

Data-points = 1.99×10^{12}

1 GPU card (2496 Cores) = 99 Minutes

1 PC (4 Cores)= 21,307 Minutes (~14.8 days)

Financial Compliance: Anti Money Laundering

Large Daily Transaction Pool
(Millions Daily)

Complex Rules
60 Expert Rules * 10 fields * 20 million customer

36 Trillion Combinations

High False Positive Rate
False positive = 99.96%

Large Financial Terrorist
(~ 6 Million Individuals)

Improved Analytical Integrity

Accelerated & Parallelized Algorithms

High Speed Search
(Seconds)

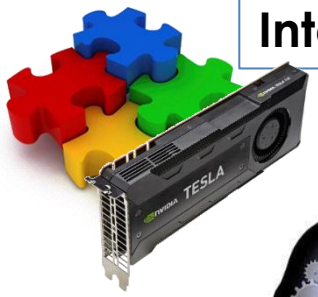

Machine Learning Algorithms
(Minutes)

Adaptive Domain Specific Algorithms
(Reduction to 20% & Hours)

Accelerated Machine Learning Solution

Daily Transactions

Bank DB

TESS DB

Historical Alerted Transactions (150K)

1. Learning from historical data

AI Trained Model

Normal Category

Suspicious Category

Predict based on Trained Model

New Daily Transactions

Bank Negara

AI Prediction Engine

Predict

Normal

Suspicious

AML Report

Report Suspicious Customers

2. Adaptive Learning based on Officer Feedback

Accuracy Test using historical data

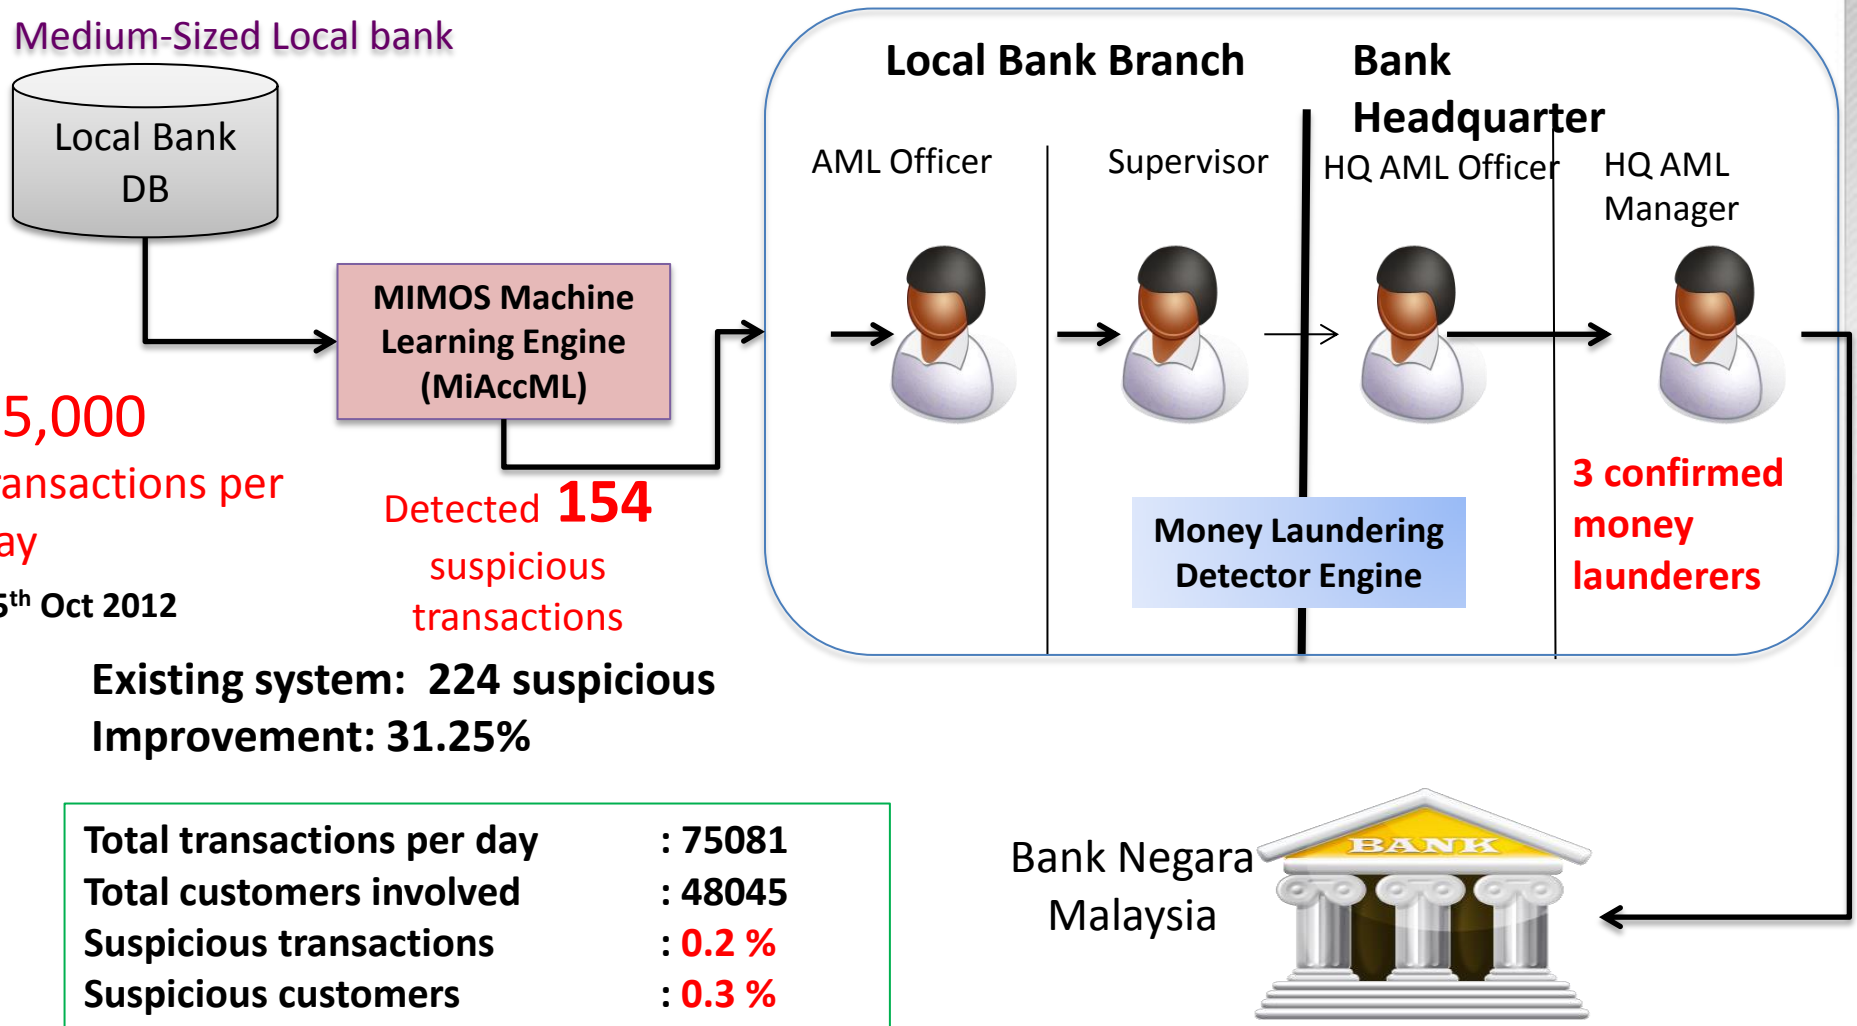

Existing system: 224 suspicious
Improvement: 31.25%

Total transactions per day	**: 75081**
Total customers involved	**: 48045**
Suspicious transactions	**: 0.2 %**
Suspicious customers	**: 0.3 %**
Actual money launderers	**: 0.009 %**

Video Analytics

Crypto

Database Acceleration

Text / string analytics

Finance

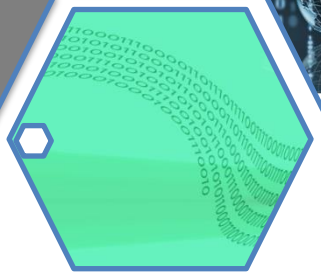

Video Analytics Implementation in GPU

* Differs based on server configuration & video complexity

GPU VA Library

Surveillance Server

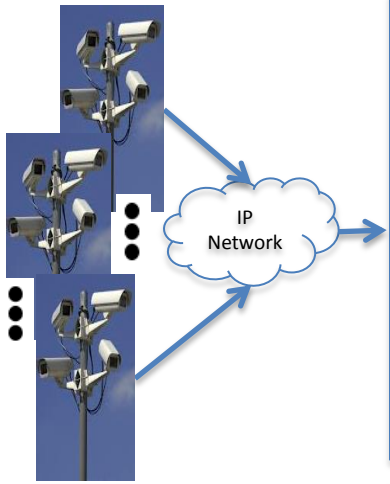

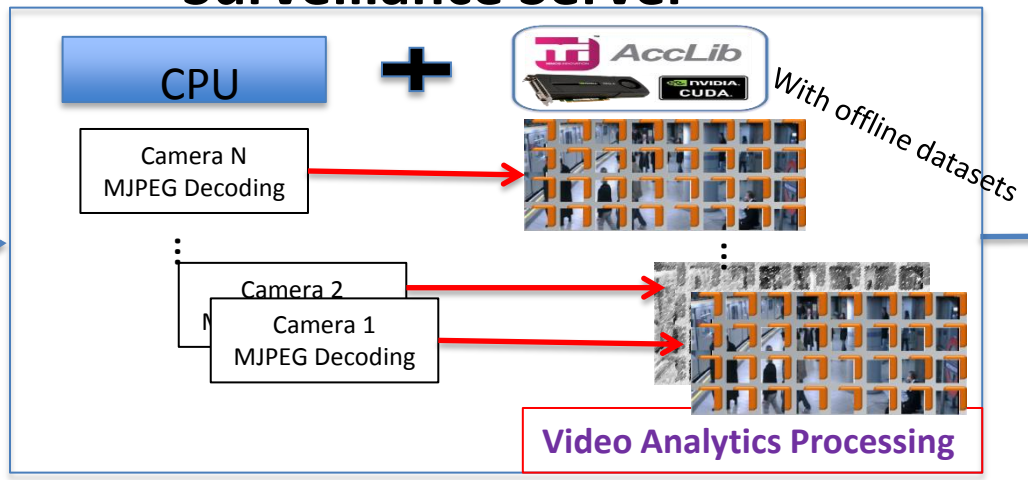

misp

Video Analytics Processing

CPU

AccLib

Parallelization of the VA algorithms

- Previous data dependency
- Efficient memory management.
- Algorithm Decomposition

CPU + GPU

Background Subtraction	AccBackgroundSubtractionFrameDiff AccCompMotion, AccUpdateBackground AccCompShadow, AccRGB2HSV
Morphing Process	AccMorphFilterVariable
CCL	AccConnectComponentLabel
Region Analyzer	AccExtractPropertiesCentroid, AccExtractPropertiesSize, AccExtractPropertiesBB, AccExtractPropertiesHWRatio, AccExtractPropertiesOrientation, AccExtractPropertiesHProject, AccExtractPropertiesSkew, AccRegionLabelUpdate, AccCompOverlap, AccPropUpdate, AccCombineBlob
Filters	AccFlickerFilter, AccRegionFilter
Detection	AccVAParallelIntrusionDetection

GPU VA System Results

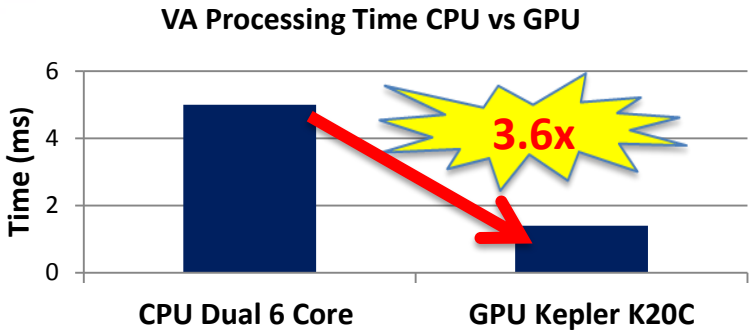

Tasks	CPU + GPU	* CPU Utilization
Network Stream In	CPU	10%
Decompression	CPU	5%
Video Analytics	GPU	35%
Streaming Out & Display	CPU	50%

* Reference to 10fps

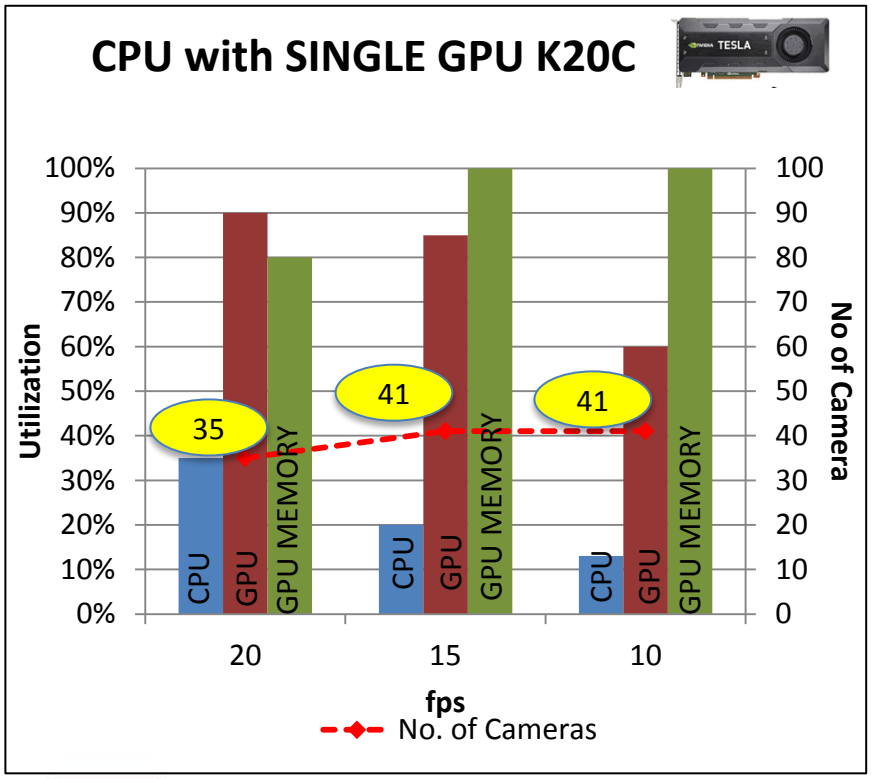

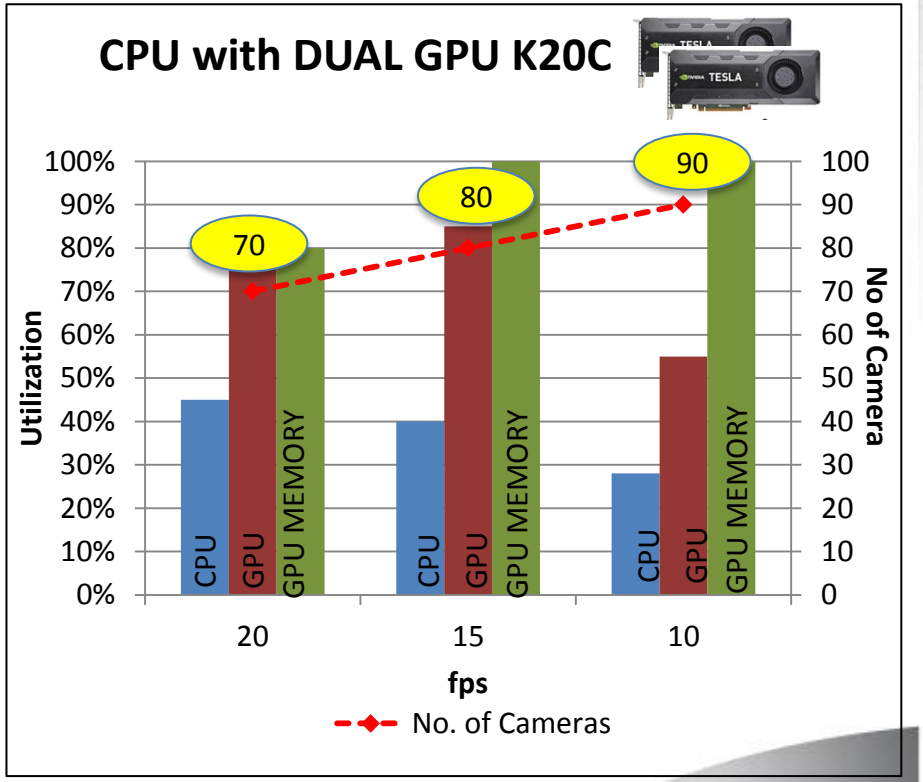

* Data taken on system server CPU - Dual 8 cores

System Savings & Roadmap

No.	Items	Setup Cost (RM '000)		Annual Maintenance (RM '000)	
		CPU 50 nodes	CPU + GPU 10 nodes	CPU 50 nodes	CPU + GPU 10 nodes
1	Space Rental	0	0	136.5	52.92
2	Infrastructure	250.5	190	15	13
3	Utilities	0	0	469.2	192
4	IT Equipment	700	780		
	Total Cost	950.5	970	620.7	257.92
	Savings		-19.5		362.78
	% Savings		-2		58

> RM300k = >50%
per annum

Crypto

Database Acceleration

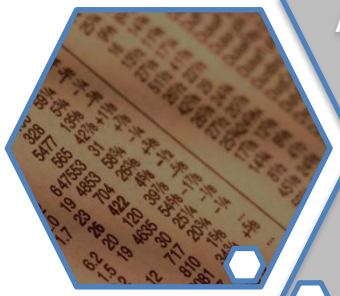

Video Analytics

Text / string analytics

Finance

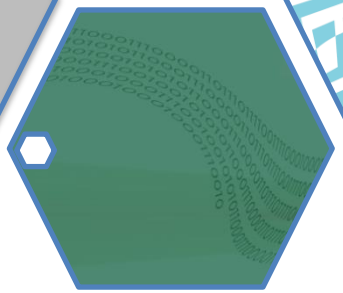

Mi-AccSSL

SSL/TLS Overview:

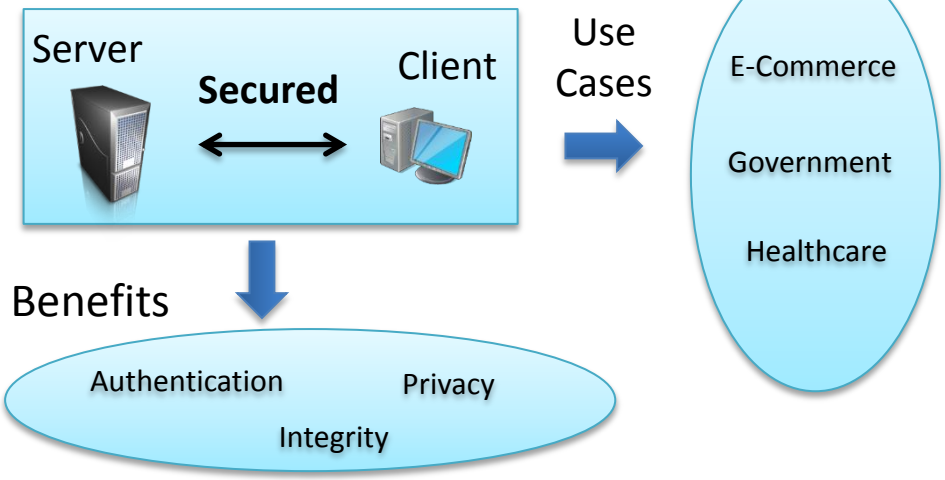

Challenges:

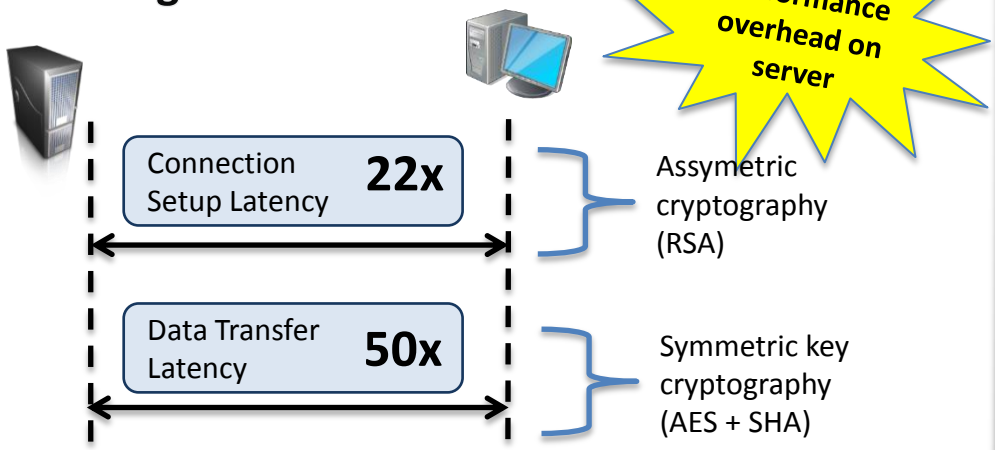

Solution:

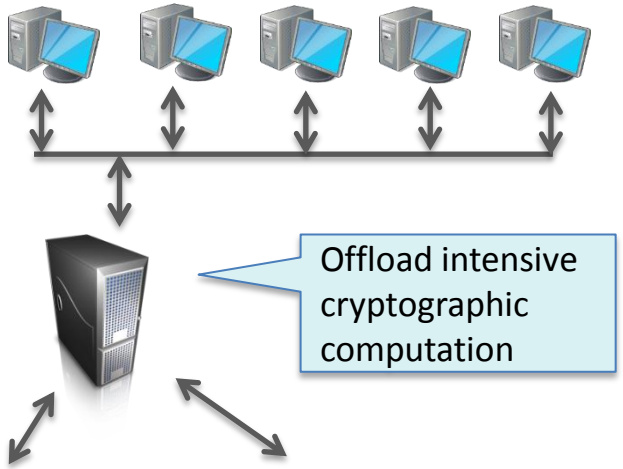

Available (A):
(Hardware accelarator)

Proposed (B):
(GPGPU based accelerator)

Criteria	A	B	Remarks
Cost		★	GPU costs **10x** lesser than hardware acc.
Flexibility		★	GPU allows addition of MIMOS crypto algorithm

AccLib Mi-AccSSL As a Platform

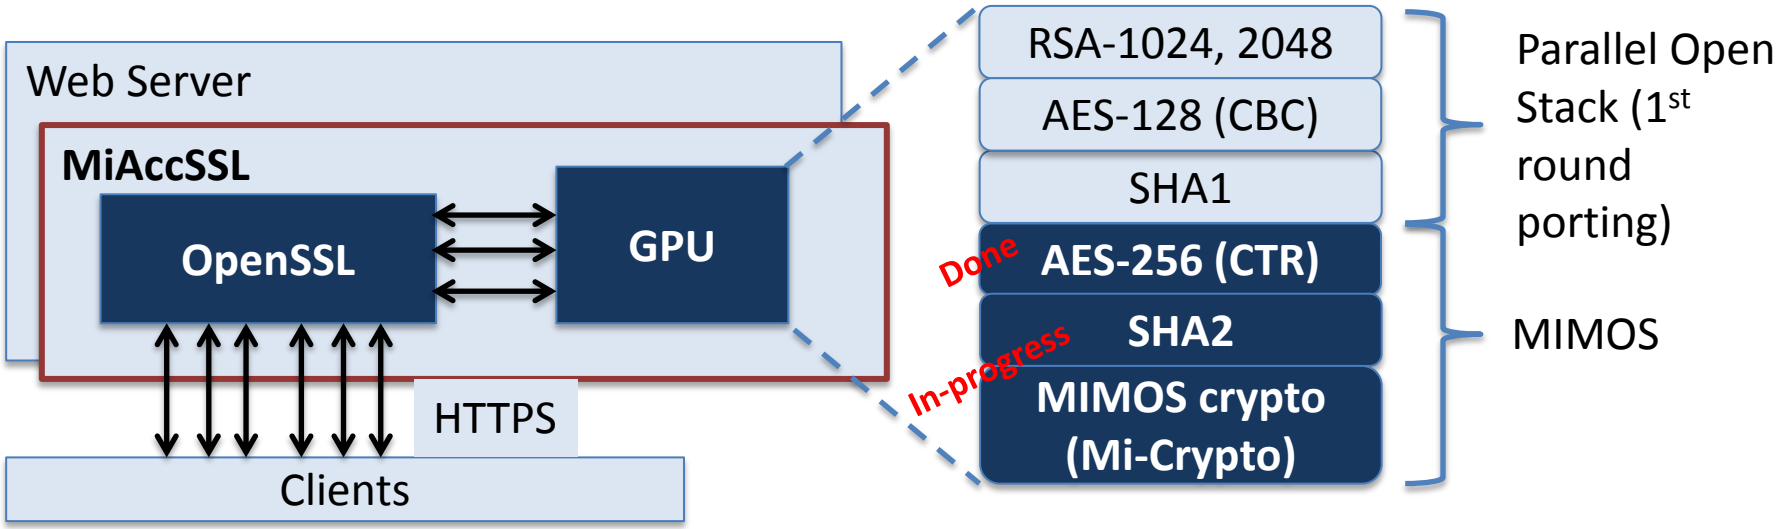

Performance:

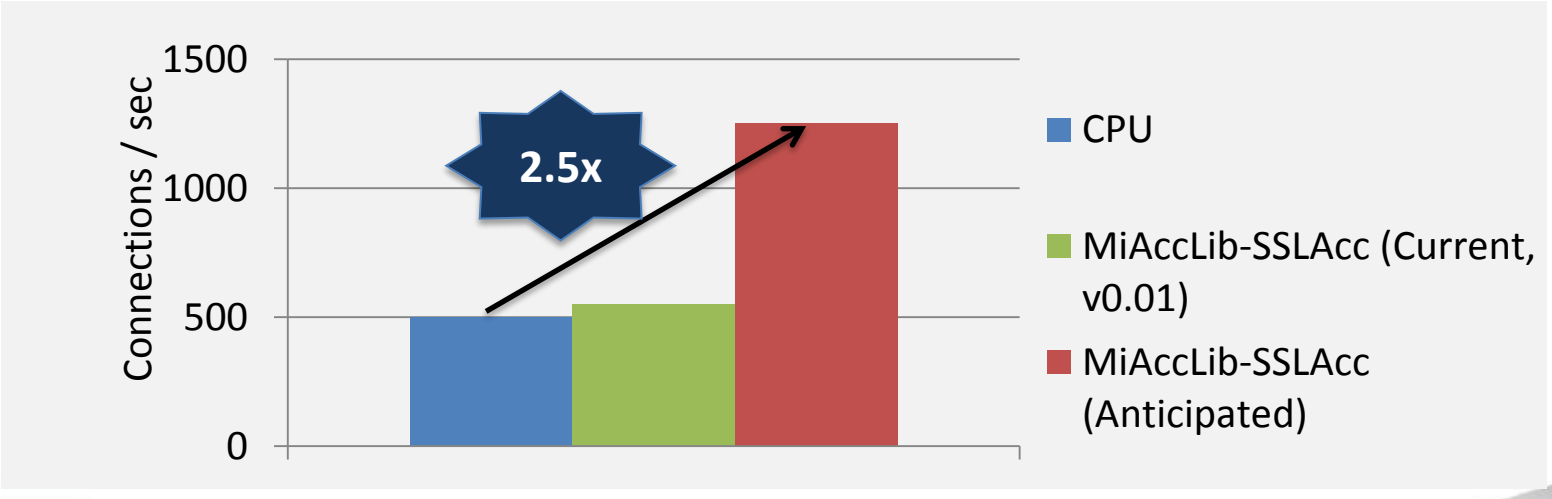

Performance comparison of **AES 128** on CPU and GPU cards

Release in MiAccLib V1.0

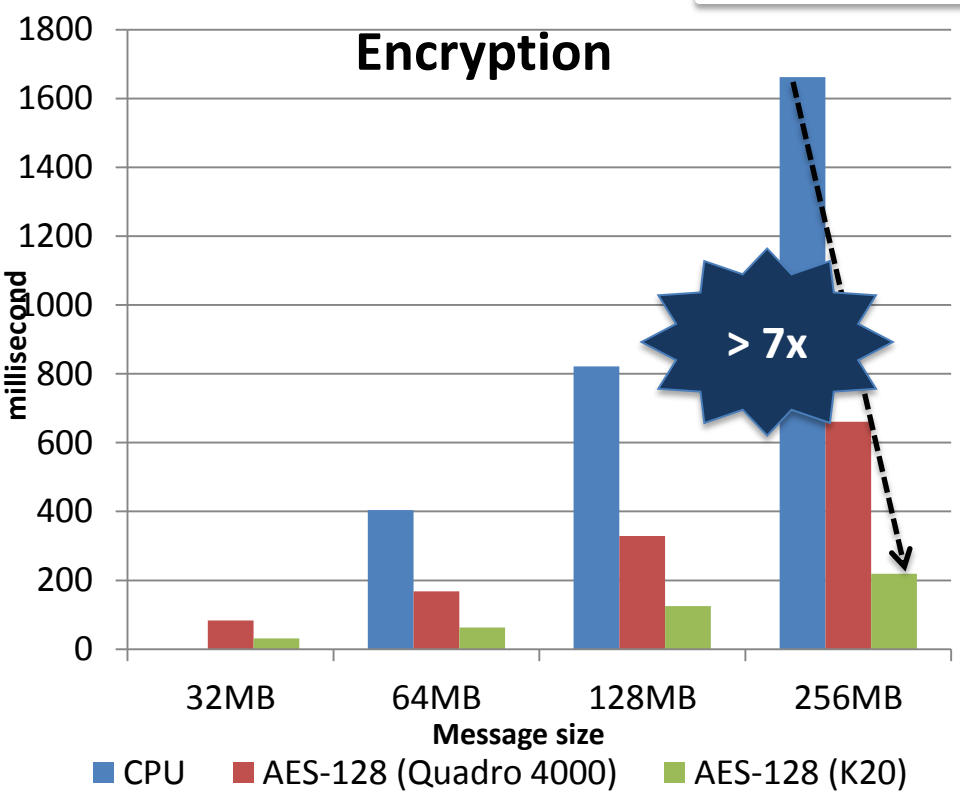

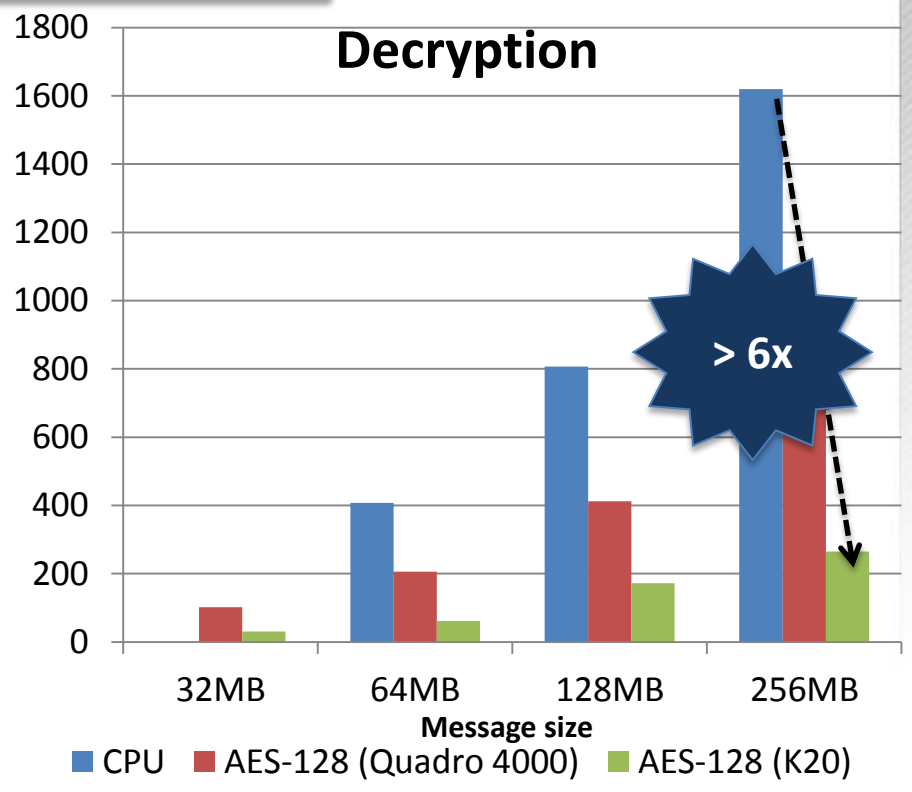

- Implemented in ECB(Electronic Code Book) mode
- ECB is easiest to parallelize but seems more vulnerable
  - Independent plaintext
  - Static key (same key)

HPC

Database Acceleration

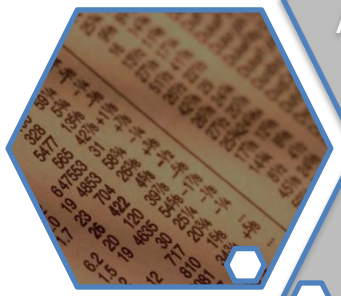

Video Analytics

Text / string analytics

Finance

High Performance Computing

Distributed machine architecture

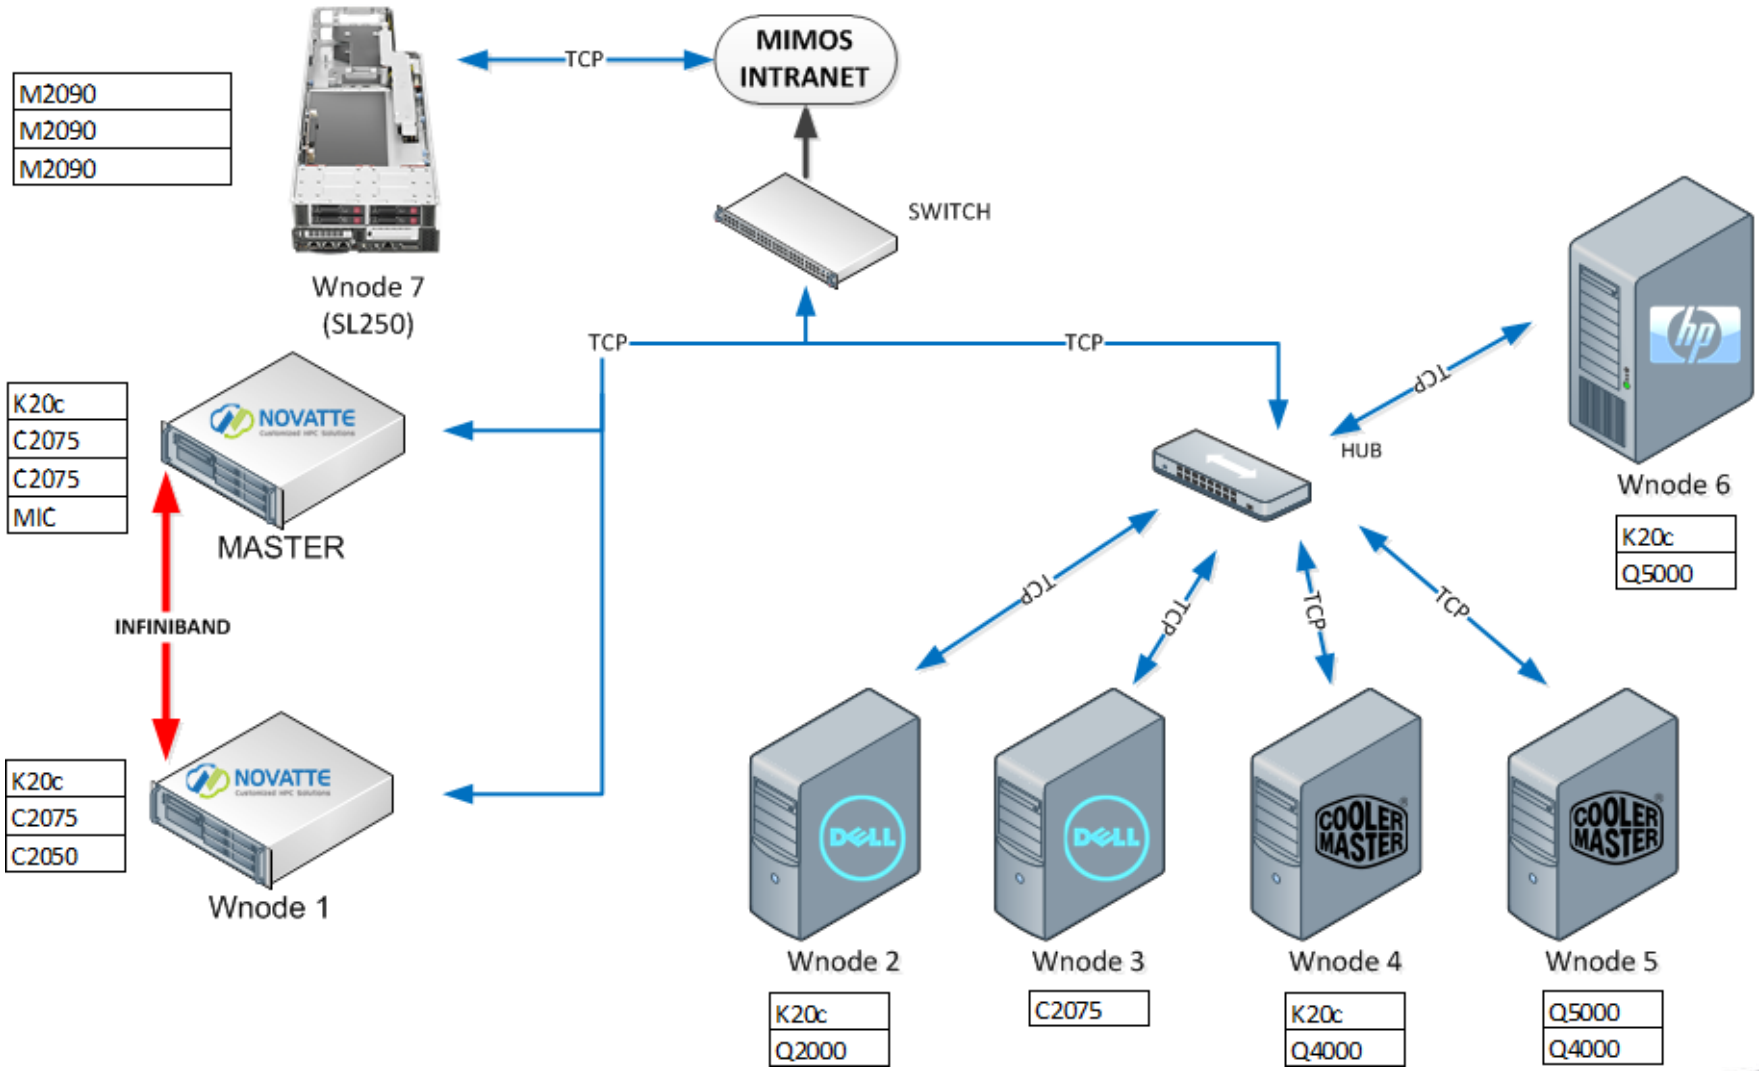

Theoretical Flops

	CPU (GFLOPS)		GPU (GFLOPS)	
master	intel Xeon E5-2640	120	K20C	3519
			C2075	1030.4
			C2075	1030.4
wnode1		120	K20C	3519
			C2075	1030.4
			C2050	1030.4
wnode2	intel Xeon E5630	40.48	K20C	3519
			Q2000	0
wnode3		40.48	C2075	1030.4
wnode4	Intel(R) Core(TM) i7 CPU 960	51.2	K20C	3519
			Q4000	486.4
wnode5		51.2	Q5000	722.304
			Q4000	486.4
wnode6	Intel(R) Xeon(R) CPU E5-2620	96	K20C	3519
			Q5000	722.304
wnode7	Intel(R) Xeon(R) CPU E5-2660	140.8	M2090	1331.2
			M2090	1331.2
			M2090	1331.2
TOTAL		660.16		29158.01
CPU + GPU (TERAFLOPS)				**29.81817**

Single Precision
29.8 Teraflops

Double Precision
~13 Teraflops

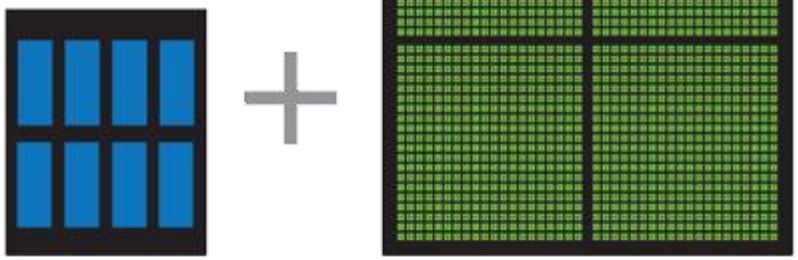

Features

- MVAPICH
- Altair PBS system (v12.0)
  - PBS Scheduling
  - PBS Display Manager
  - PBS Compute Manager
  - PBS Analytic
- NFS
- MIC / GPU / MULTICORE
- Point To Point Mellanox Infiniband
- GPU Direct

Altair | PBS Works™

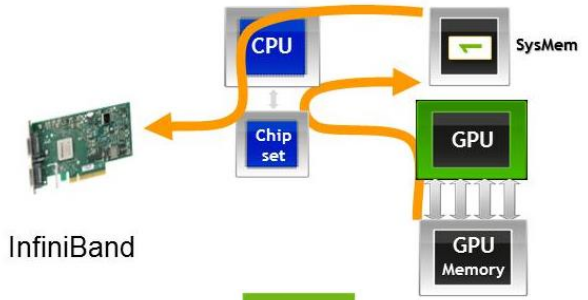

NVIDIA®

MVAPICH

GPU R&D and Compute Solution Center

- Collaboration
- MiAccLib promotion
- Training
- GPU workshop
- Roadshow / Conference
- TAP program

GPU R&D and Compute Solution Center

- First R&D Eco-system for GPU (Parallel Computing) in South East Asia
- Launched in 18 Oct 2012

MiAccLib – Promoted by NVIDIA

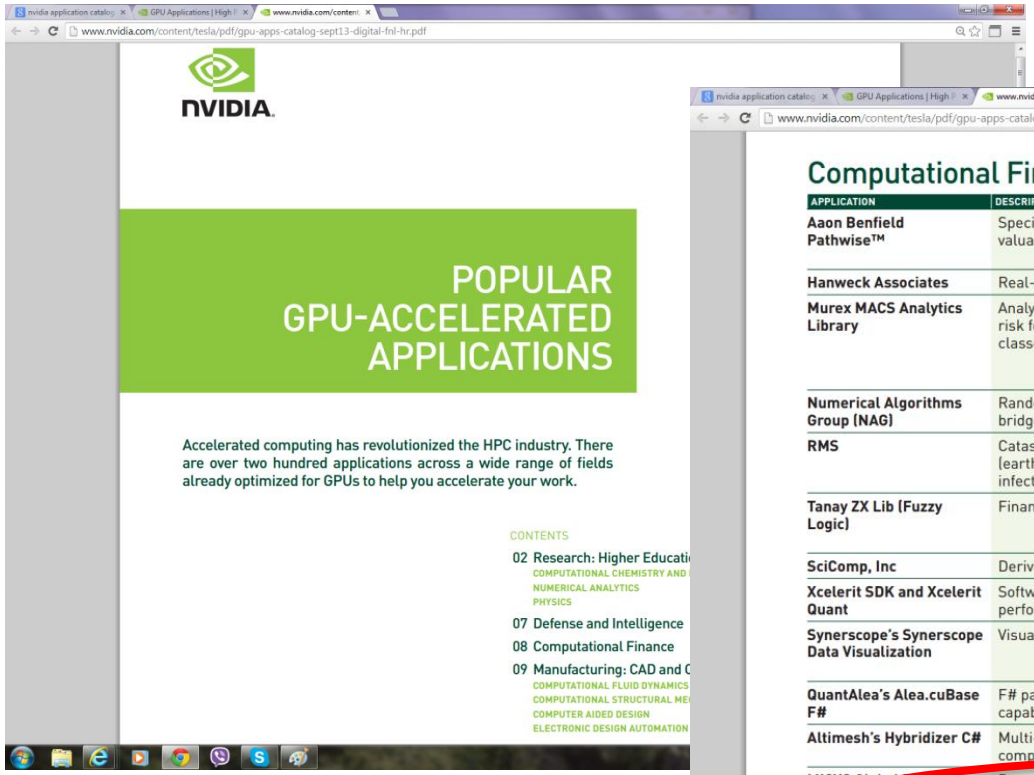

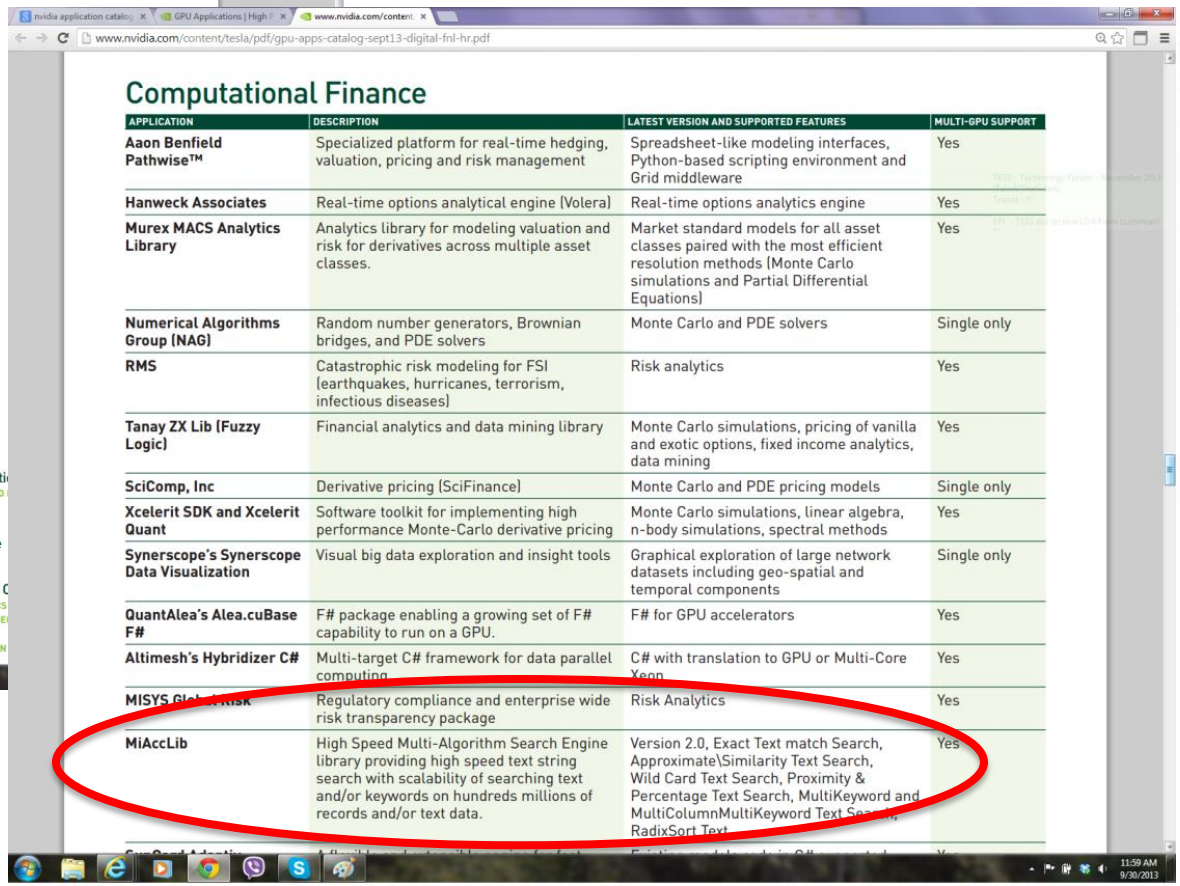

Release 1.0
(new functionalities under development & verification)

http://www.nvidia.com/docs/IO/123576/nv-applications-catalog-lowres.pdf

Showcases of GPU technologies and leading GPU enabled applications

GPU R&D AND COMPUTE SOLUTION CENTER

by ACCELERATIVE TECHNOLOGY LAB

http://gpu.mimos.my

Sponsors

About GPU R&D AND COMPUTE SOLUTION CENTER

The first in South East Asia, GPU Solution Center in Malaysia is a unique collaborative environment between NVIDIA, HP, MIMOS and key software vendors, that serves to display and share the latest GPU technologies together with GPU enabled applications such as SolidWorks, Autodesk 3ds Max, ANSYS, Abaqus, Adobe CS6, and others.

Besides latest GPU application showcasing, this center also caters for various GPU related R&D activities such as GPU Accelerator Libraries Development, GPGPU enabled Application Specific/Generic Algorithm R&D, GPGPU application benchmark & testing.

The Center helps promote the dynamic growth of GPU Computing in Malaysia through collaborative university research projects, CUDA teaching center formation at universities, annual GPGPU workshop, annual CUDA programming contest, technology road shows, various Government & private entity visits hosting and other GPU related activities.

GPU R&D and Compute Solution Center Objectives:

- To facilitate adaptation of many-core/parallel/GPU techniques in scientific, financial, big data processing areas
- To enhance GPU related R&D activities in Malaysia
- To serve as a one-stop center to promote, share & teach GPU technologies/solutions to customers and those interested in GPGPU, and to do joint collaborations on GPU topics

Vision

- To set GPU R&D and Compute Solution Centre as a Center of Excellence for pioneering work in parallel computing research in South East Asia using accelerated technology and GPUs utilizing integrated compute solution platforms
- To embark on research projects and education activities for the continued interest and benefits in massively parallel computing and NVIDIA, HP, and MIMOS technology
- To promote interests in taking GPU-based solutions to indigenous local industries

Mission

- To conduct research and development on the use of GPU personal supercomputing platforms as well as GPU-enhanced multi/many-cores platforms
- To engage in a number of research, development and educational activities (e.g., applications, software development tools, system software and architectures) that leverage heterogeneous & scalable computing
- To provide expertise, information, guidance and tools for affiliate members to engage in GPU & Compute-related projects
- To be at the forefront of computing technology for developing state-of-the-art computational algorithms that drive innovation in the industrial application, sciences and engineering
- To promote participation of Affiliated members to be involved in GPU-related activities by cross-fertilizing ideas and skills, by sharing software and hardware facilities, leveraging training materials and efforts, and by streamlining interactions with closer, priority access to NVIDIA/HP staff and capabilities

Promote collaboration with university

UTM, MMU & USM (2013); Uniten & USM (2014

Annual GPU Workshops

- GPU workshops organized to date @MIMOS
  - GPU Workshop 2011
  - GPU Workshop 2012
  - GPU Workshop 2013

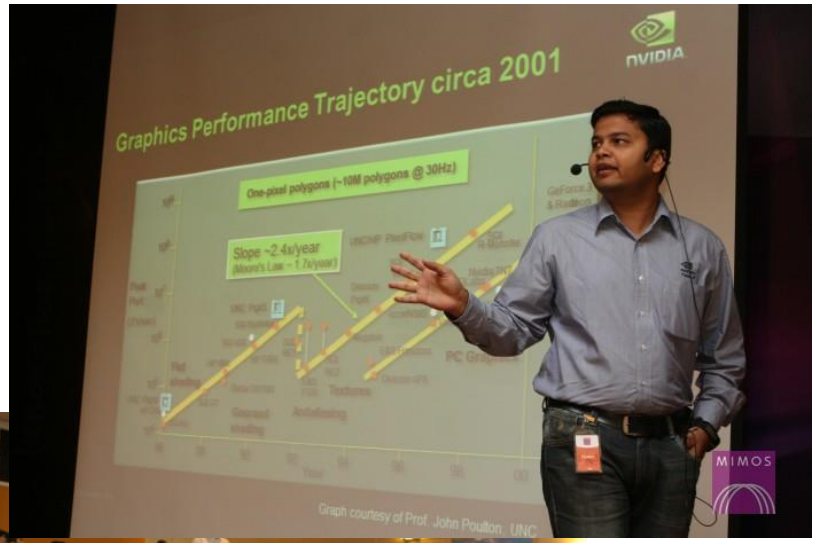

Technology Road Shows & Awareness

- Support Technology Road Shows targeting specific verticals
  - Oil & Gas (February 2013)
  - Big Data (2013)
  - Finance (2012 + 2013)
  - M&E (2013)
- Promote Technology Innovation

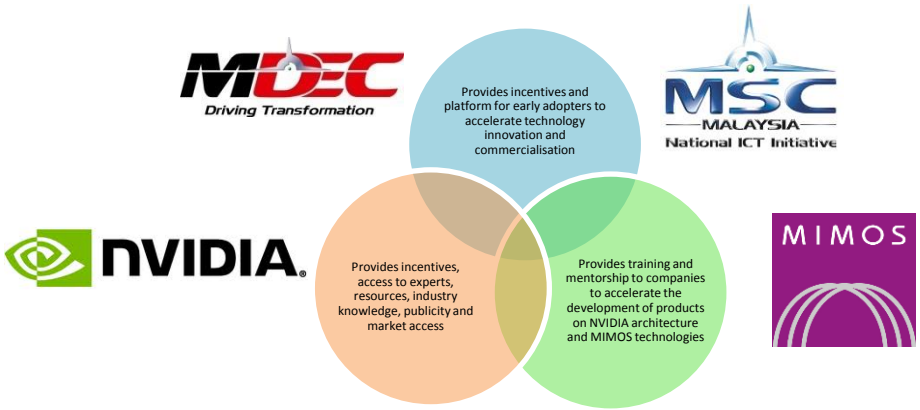

TAP Program – Leveraging Network

Pending Term-sheet

TLA Done, Integrated & Testing In Progress

Pending Term Sheet

TERIMA KASIH

THANK YOU

www.mimos.my

Innovation for Life™