A Sketch of Data (graph) Analytic Applications in the Medical Field

with Thoughts on the Applicability of CnC as a Framework for Hybrid Platforms in These Application Spaces

Gary S. Delp, PhD

Just a simple engineer

Mayo Clinic

Special **P**urpose **P**rocessor **D**evelopment **G**roup

Concurrent Collections (CnC) Workshop

7 September 2015

MAYO CLINIC

Purpose and Agenda

- Collect your wisdom on: ***the Applicability of CnC as a Framework for SPPDG-class Hybrid Platforms***

Agenda

- **BACKGROUND AND TERMINOLOGY**
  - **Dwarfs, Hybrid Platforms, Elfs (the plural spelling is intentional)**
  - **Constraints, Algorithms, Tuning Specifications, Tags and Collections**
  - **CnC: the abstraction is not just the implementation**
- **MEDICAL APPLICATIONS & A DWARFS USAGE GRID**
- **COMPARATIVE PERFORMANCE DWARFS AND PLATFORMS**
- **HIERARCHICAL CNC**
  - **The Elfs manage the Dwarfs, they all obey their domain constraints**
- **STREAMING ANALYTICS**
- **CNC APPLICABILITY: THE GOOD, THE FUTURE, and THE CONFUSING**

Background

- The Special Purpose Processor Development Group (SPPDG), one of the many research labs at the Mayo Clinic in Rochester Minnesota, has been studying applications that are served well by a variety of processing architectures. These include NUMA vector processors, single-threaded, high-speed processors, GPUs, FPGAs and branch optimized processors.
- We present sketches of example applications and architectures that have a variety of "impedance matching" (problem to platform) characteristics.
- In the "Big Data" field, these problems are not the Giants (big but well formed), but rather the Ogres (Fox et al.).

Terminology

Some of the terms used in this report take on specialized meanings. These terms are used with these specialized meanings throughout the talk.

Elfs

- Elfs will be used to attack the Ogre-shaped problems.
- Processing elves, unlike the Berkeley Dwarfs (Asanoviç et al.) are combinations of low-level dwarfs with a high-level (distributed) view of the data.
- Elfs are powerful, and can be used repeatedly. Elfs are long-lived and close to tireless.
- Different from specific workflows, one elf can be used in many workflows.
- Elf-based results can illustrate the utility of- and the need for- considering a very large number of factors with potentially subtle interactions.

The Elf and Dwarf Dependencies

- The inter-processor communication of events and transport of data make processor affinity often more important than processor architecture.
- Work on exploring frameworks that can work across and between these islands of capability is ongoing.
- This exploration has indicated needs for dynamic affinity scheduling, low cost nonce value abstraction (running out of nonce identifiers, or having them centrally managed is a potential issue), and SQL and graph database interactions
- Streaming applications address locality limitations in time and space; the various data structure & storage architectures are attempts to find long baseline correlations.

High Performance Data Analytics (HPDA)

- Generally, *data analytics* is a loosely used term. It is often used interchangeably with *graph analytics* or *big data*. In this talk, high performance data analytics (HPDA) involves analyzing enormous data sets with complex *non-regular* relationships to discern patterns that are extremely non-local.
- The non-locality and irregularity of the relational data require the need for any processor/thread to be able to access any portion of the entire (huge) data set. This increases the computational challenges significantly.
- A canonical example is the analysis of Facebook users and their *friend* relationships, represented as complex graphs (users =nodes, relationships=edges, with additional data, such as duration, timestamps, etc. represented as alternate types of edges or nodes).
- A large-scale computing counter-example to HPDA would be a massively *embarrassingly parallel* computation (*e.g.*, a Monte Carlo simulation of light transport between insertion and detection through a complex medium) in which very large *aggregate state* is be held and processed at one time, but each processing element needs access to a small (traditionally cacheable) amount of this total state

Hybrid Computing Platforms (HCP)

- The bulk of existing systems that are currently referred to as hybrid computing systems include more than one processor type, *e.g.*, CPU & GPU, and require programmatic block transport of data between the computing units. If memory space exists – that is shared between and amongst the various processors – it is limited.
- As used in this talk, a Hybrid Computing Platform (HCP) ideally contains
  - Globally accessible but physically distributed memory
  - hardware supported thread migration
  - multiform processors
  - memory side processing, including widespread and selectable in-memory synchronization. These features are not currently available in commodity hardware.

The Computational Giants of Massive Data Analysis (Adopted from [1])

Committee on the Analysis of Massive Data
Committee on Applied and Theoretical Statistics
Board on Mathematical Sciences and Their Applications
Division on Engineering and Physical Sciences
National Research Council of The National Academies

Giant	Name
G1	Basic Statistics
G2	Generalized N-Body Problems
G3	Graph-Theoretic Computations
G4	Linear Algebraic Computations
G5	Optimizations
G6	Integration
G7	Alignment Problems

[1] National Research Council, Frontiers in Massive Data Analysis, Washington, DC: The National Academies Press, 2013. Available: http://www.nap.edu/catalog/18374/frontiers-in-massive-data-analysis.

THE VIEWS THAT CAN BE TAKEN OF BIG DATA OGRES AND THE FACETS OF THOSE VIEWS (The *Views* Include Data Source and Style, the Problem Architecture, Execution, and the Processing View; This Is From Early Work By Fox, et al., On Developing A Systematic Approach To *Big Data* Benchmarking)

Ogres

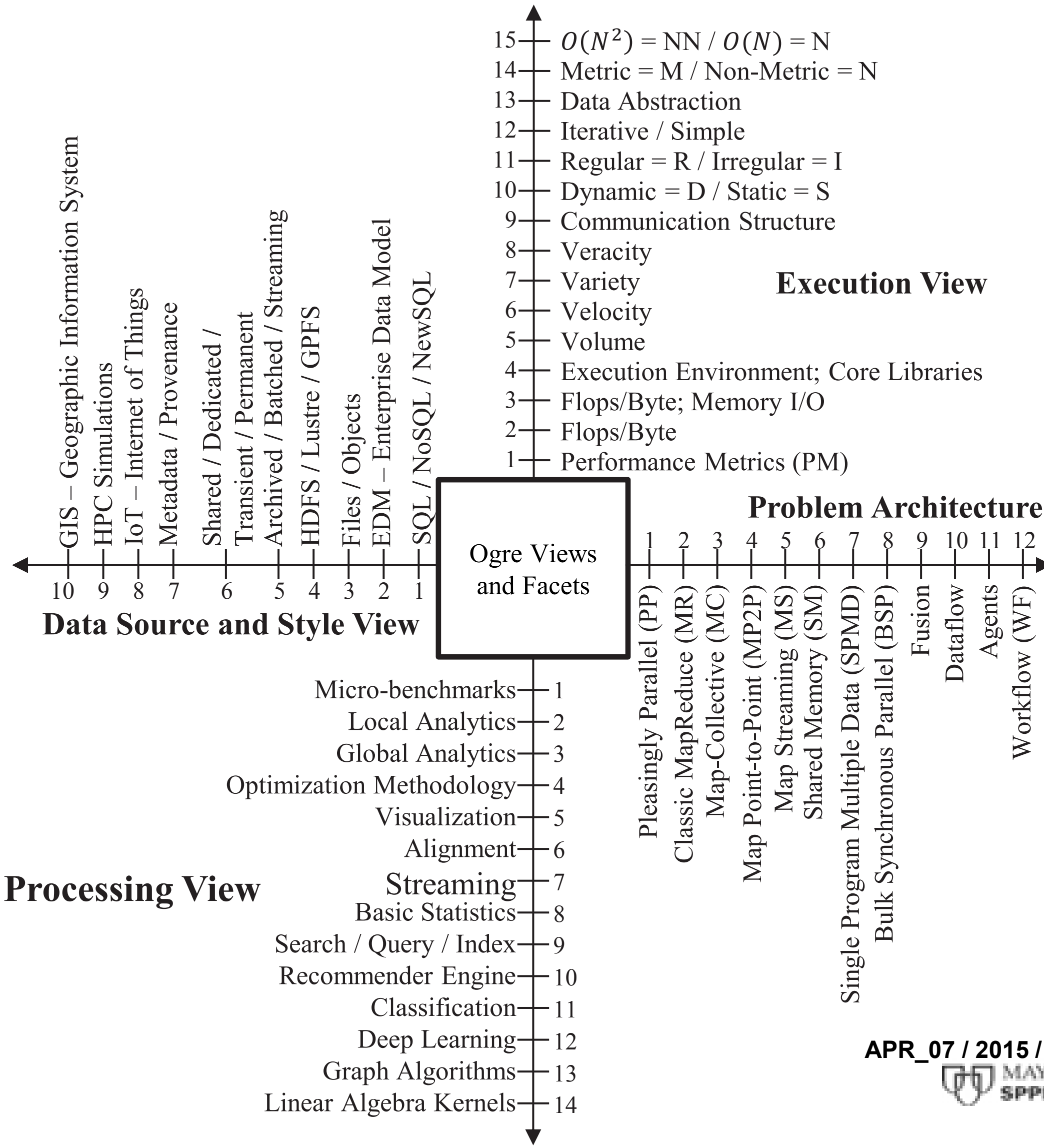

APR_07 / 2015 / GSD / 44838

MAYO CLINIC SPPDG

Adapted from, Fox, G.C., et al.: "Towards a Systematic Approach to Big Data Benchmarking," Community Grids Lab: Pervasive Technology Labs, Computer Science and Informatics, Indiana University, Bloomington, IN, Technical Report submitted for publication, 15 February 2015; http://grids.ucs.indiana.edu/ptliupages/publications/OgreFacetsv9.pdf

Dwarfs

- Phil Colella is credited with the recognition of the *Seven Dwarfs* in his 2004 presentation “Defining Software Requirements for Scientific Computing” about DARPA’s High Productivity Computing Systems (HPCS) program [3]. Berkeley’s View project [1] added to the list of dwarfs, keeping the spelling used by Colella.

The Dwarfs of Berkeley

- These dwarfs are classes of structured algorithms. Abstracted from the Berkeley report, they classify algorithms (or sub-algorithms) that are similarly characterized by memory access patterns, scalability, computation intensity, mix of operations, etc. SPPDG directly adopted, and expanded some of the dwarfs to use as column headings for the low-level algorithms in Table 1.

The Original Seven Dwarfs

- **Dense Linear Algebra** (e.g., BLAS [Blackford et al 2002][13], ScaLAPACK [Blackford et al1996] [14], or MATLAB [MathWorks 2006], MATLAB[15])
- **Sparse Linear Algebra** gains it sparse name because the data sets include many zero values. Data are usually stored in compressed matrices to reduce the storage and the bandwidth required to access the remaining, *nonzero* values.
- **Spectral Methods** (e.g., FFT [Cooley and Tukey 1965][16])
- **N-Body Methods** depend on interactions between many discrete points.
- **Structured Grids** in which data are represented by a regular grid; points on grid are conceptually updated together; it has high spatial locality.
- **Unstructured Grids** comprise an irregular grid where data locations are selected, usually by underlying characteristics of the application. Data point location and connectivity of neighboring points must be explicit.
- **Monte Carlo** / later expanded to **MapReduce** in which calculations depend on statistical results of repeated random trials. This dwarf is considered embarrassingly parallel.

Two Machine Learning Dwarfs

- **Dynamic programming** is an algorithmic technique that computes solutions by solving simpler overlapping sub problems. It is particularly applicable for optimization problems where the optimal result for a problem is built up from the optimal result for the subproblems.
- **Backtrack and Branch-and-Bound**: These involve solving various search and global optimization problems for intractably large spaces. Some implicit method is required in order to rule out regions of the search space that contain no interesting solutions. *Branch-and-bound* algorithms work by the divide and conquer principle: the search space is subdivided into smaller subregions ("branching"), and bounds are found on all the solutions contained in each subregion under consideration.
- **Neuromorphic Computing**: computing that is shaped like what we think the brain does with its neurons. Picture next page

NVidia: \\millenium\Reference\ClassMaterial\2015\NVidia_DeepLearning

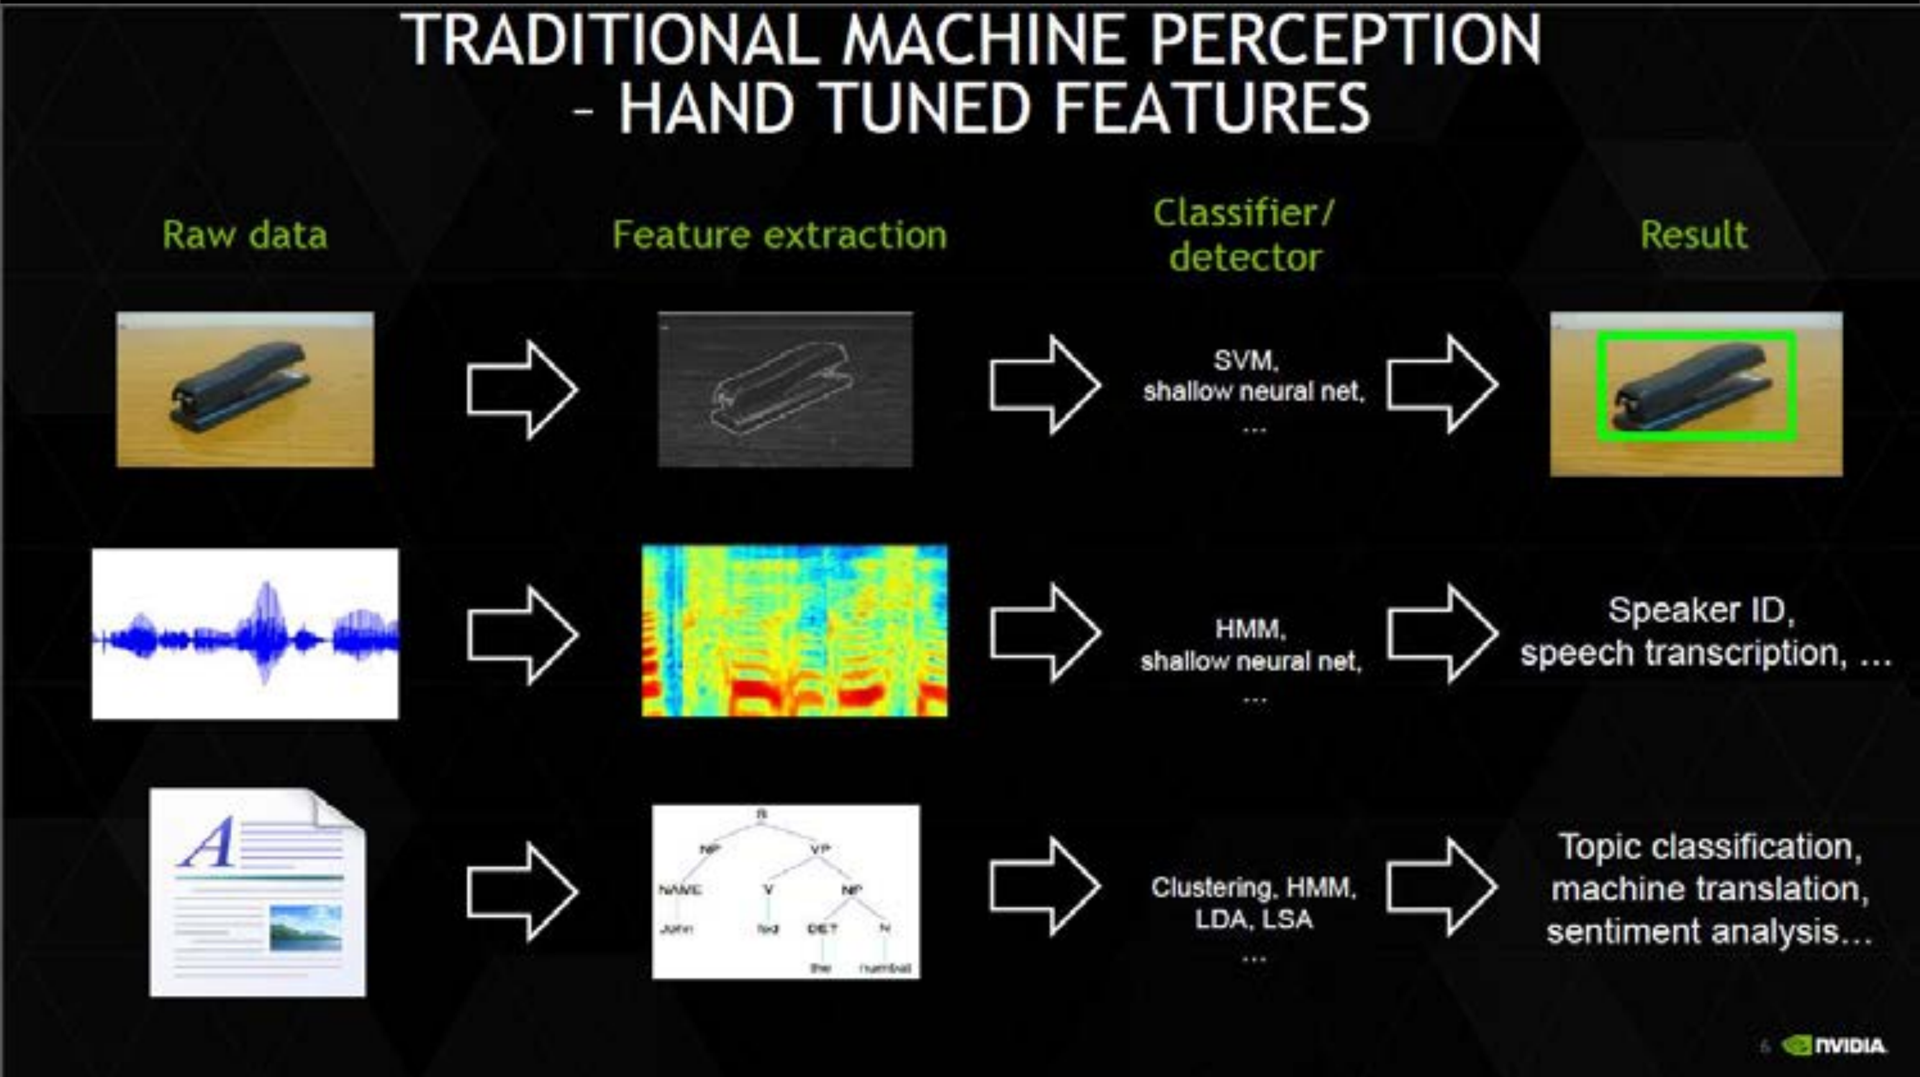

NVidia: \\millenium\Reference\ClassMaterial\2015\NVidia_DeepLearning

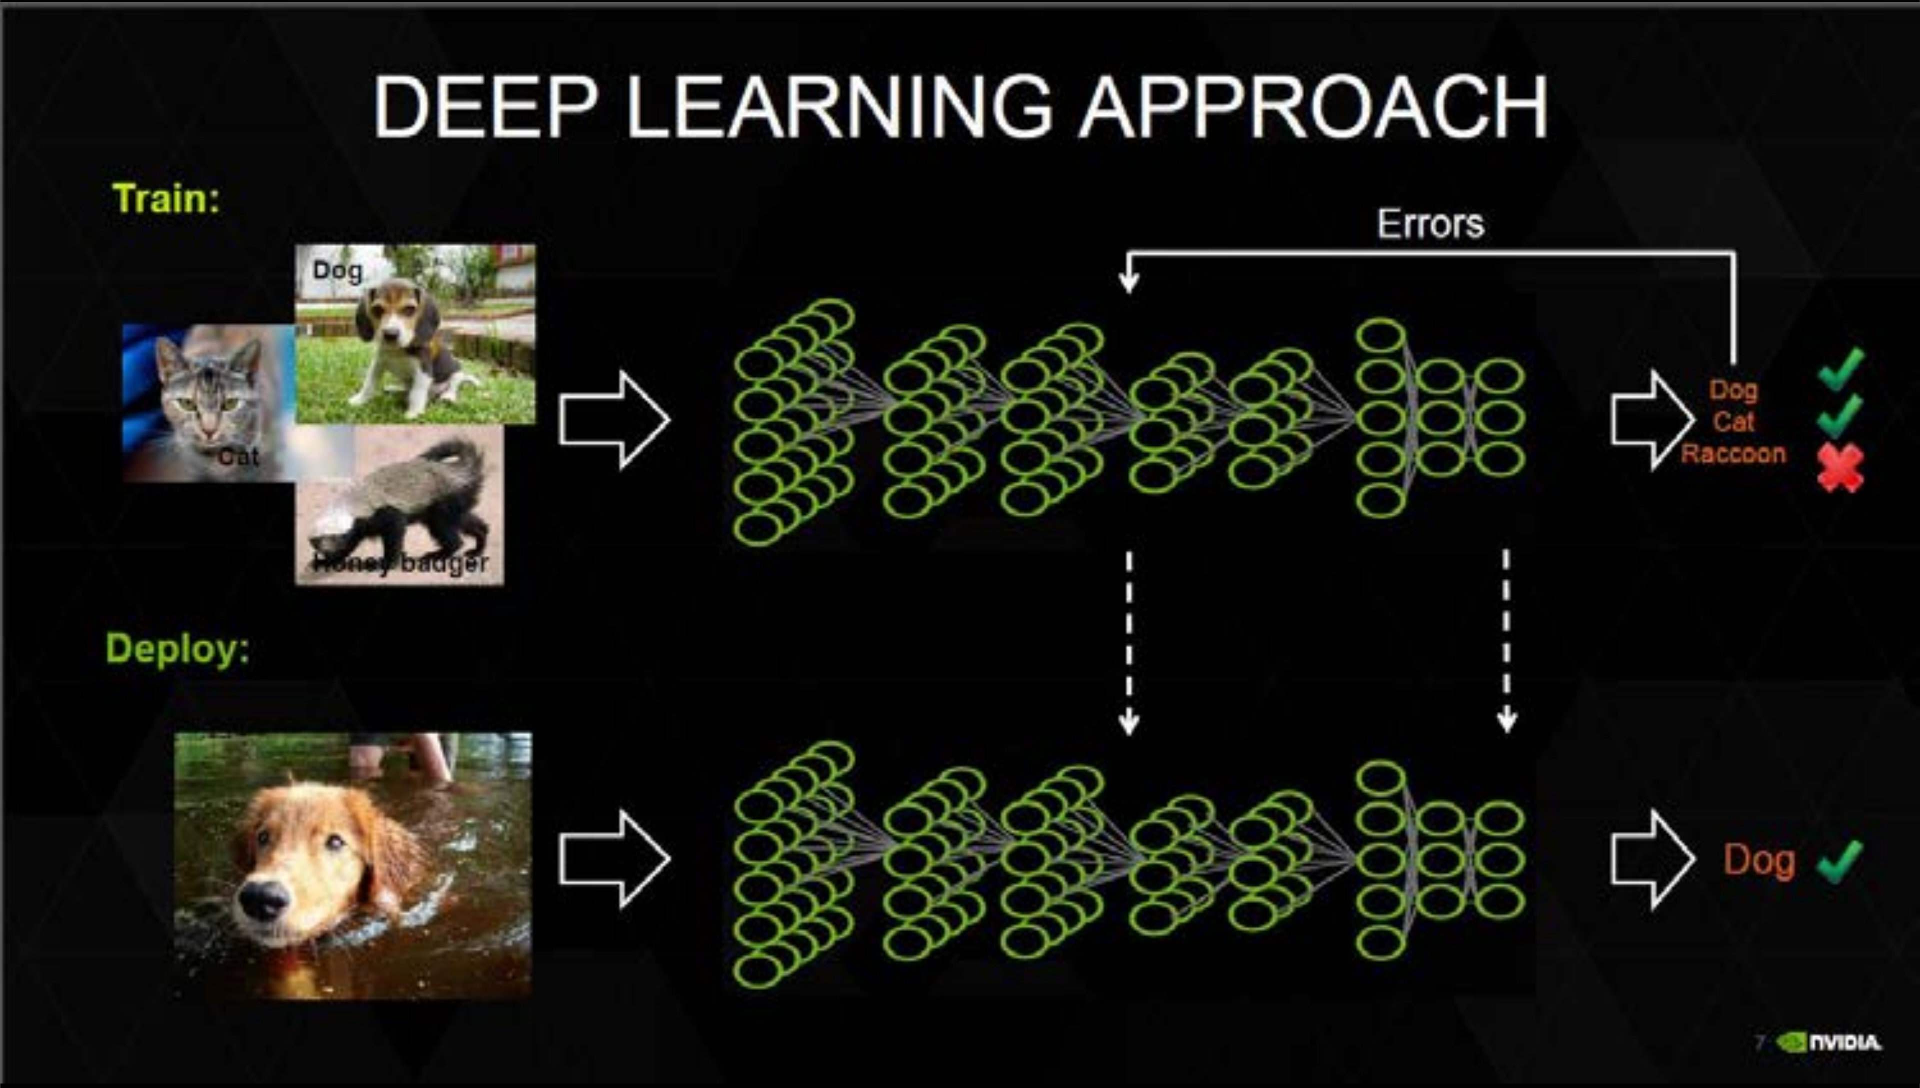

Four More Dwarfs

- **Combinational Logic** generally involves performing simple operations on very large amounts of data often exploiting bit-level parallelism. For example, computing Cyclic Redundancy Codes (CRC) is critical to ensure integrity and RSA encryption for data security.
- **Graph Traversal** applications must traverse a number of objects and examine characteristics of those objects such as would be used for search. It typically involves indirect table lookups and little computation. [*The* graph traversal *dwarf has grown to an army of giants (elfs); it is a field of important and active research and the development of* graph traversal *elfs will provide fodder for new architectural optimization.*]
- **Graphical Models** applications involve graphs that represent random variables as nodes and conditional dependencies as edges. Examples include Bayesian networks and Hidden Markov Models. [*Compressive sensing and principal component analysis are two key new subsets of this dwarf. In Table 1, this dwarf has been combined with the previous dwarf (11).*]
- **Finite State Machines** represent an interconnected set of states, such as would be used for parsing. Some state machines can decompose into multiple simultaneously active state machines that can act in parallel.

Elfs

- Not the oxymoron it seems, the *high-level* dwarfs referenced in the presentation [2] are the combination of low-level dwarfs with a high-level view of the data.
- We coined the term Elf (and the attendant plural misspelling Elfs following [3]) as a convenient label. Although, not used in the literature, it is starting to be adopted by some.
- Elfs take a high-level view. They are powerful, and can be used repeatedly. Elfs are long-lived and close to tireless. Different from specific workflows, one elf can be used in many workflows. Several of these elfs illustrate the utility of- and the need for- considering a very large number of factors with potentially subtle interactions.
- The existence of Elfs support

  Many-Factor, Subtle-Interaction Data Analytics.
- The result of two relevant Google searches return 'No results found for big-data dwarfs ogres giants analytics +elfs' and 'No results found for big-data dwarfs ogres giants analytics +elves' [4: Delp, 2015].

Patient Cohort

- *A Patient cohort* is created, possibly in real-time, from a collection of similar and related data records to be used in evaluating a *focus patient*.
- Previously, the creation of a *patient cohort* specific to a *focus patient* has been computationally prohibitive. The databases necessary to refine a *cohort* for comparison and decision support has not been available.
- With the existence of the Mayo enterprise data trust (EDT) [8], and the wide-speed, mandated development of the electronic health record (EHR), alongside the availability of advanced hybrid computational platforms and algorithms, the development of medically significant *patient cohorts* becomes potentially feasible. Although some have discussed this concept, it is not, in the knowledge of the authors, in general use.

FINDING THE MISSING LINK FOR BIG BIOMEDICAL DATA

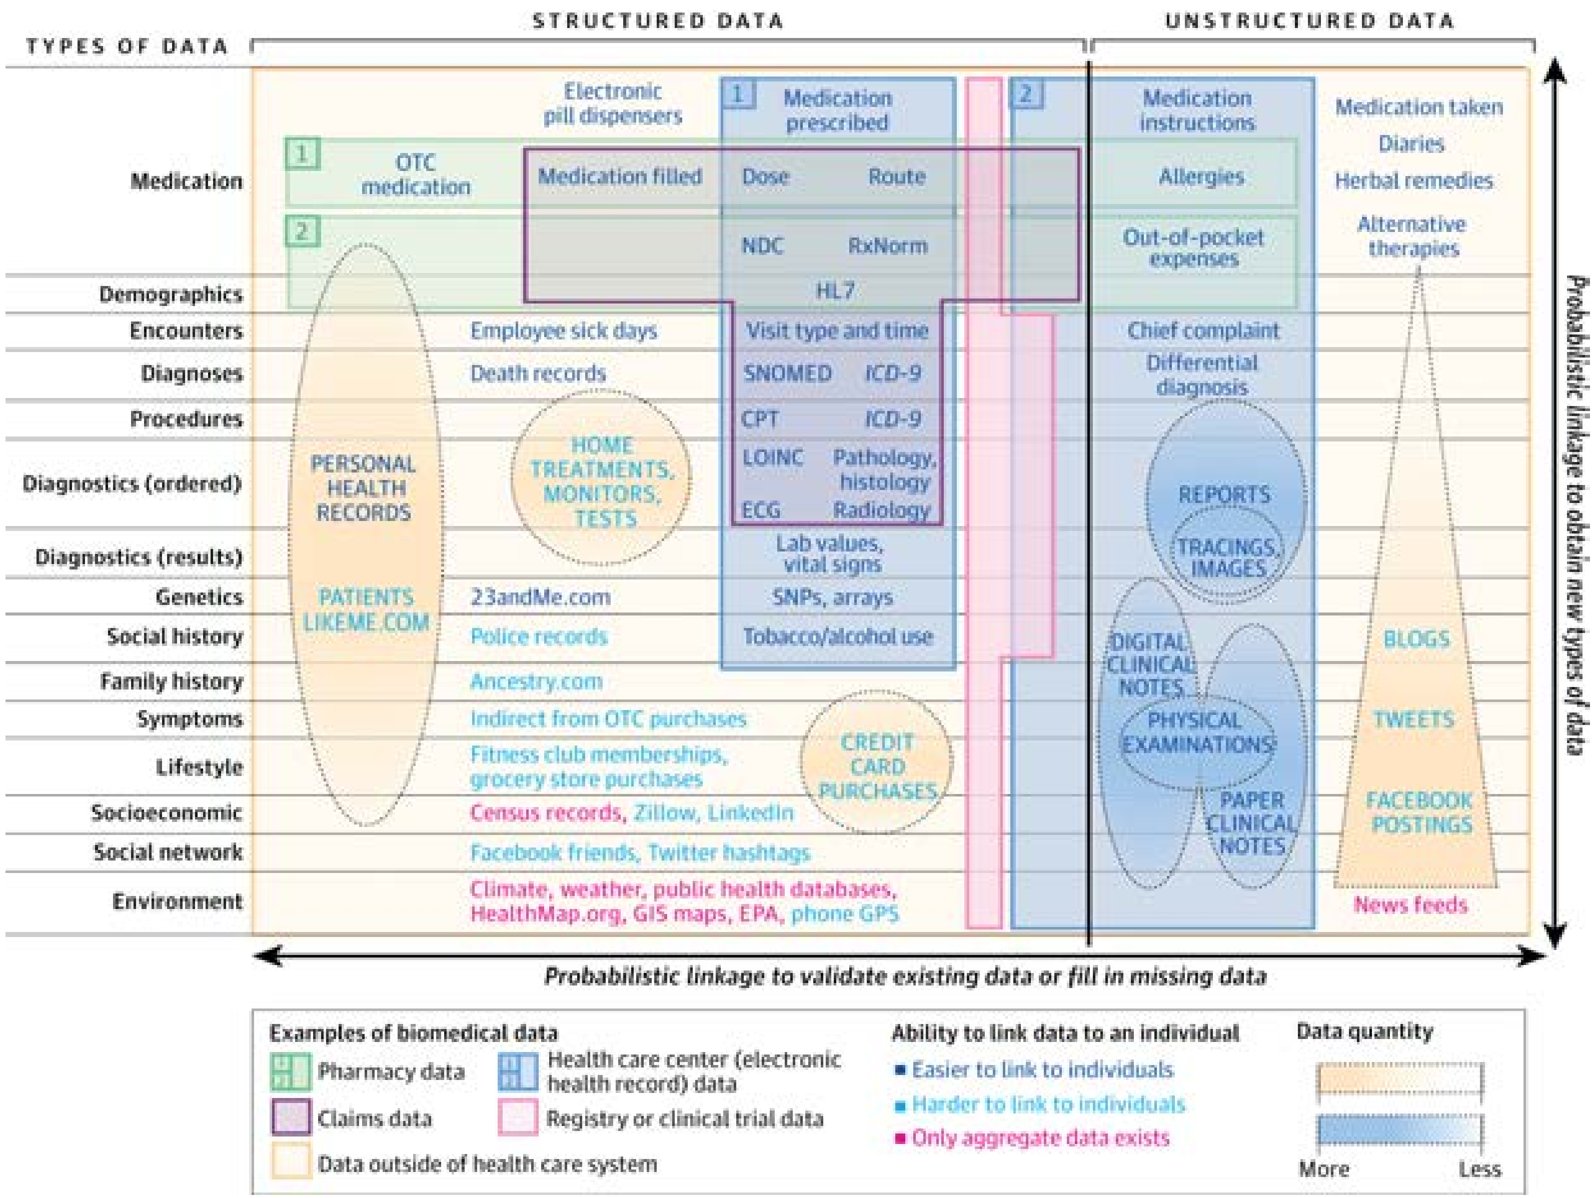

Used With Permission from, Weber, G.M., K.D. Mandl and I.S. Kohane: "Finding the Missing Link for Big Biomedical Data", JAMA, 311(24):2479-2480, 2014; http://jama.jamanetwork.com/article.aspx?articleid=1883026.

WHY DO WE NEED TO SUPPORT PREDICTIVE ANALYTICS IN MEDICINE? PREDICTIVE MEDICINE AS THE NEW PHILOSOPHY IN HEALTH CARE

Contrasting Healthcare Philosophies (Interventional vs. Preventative)

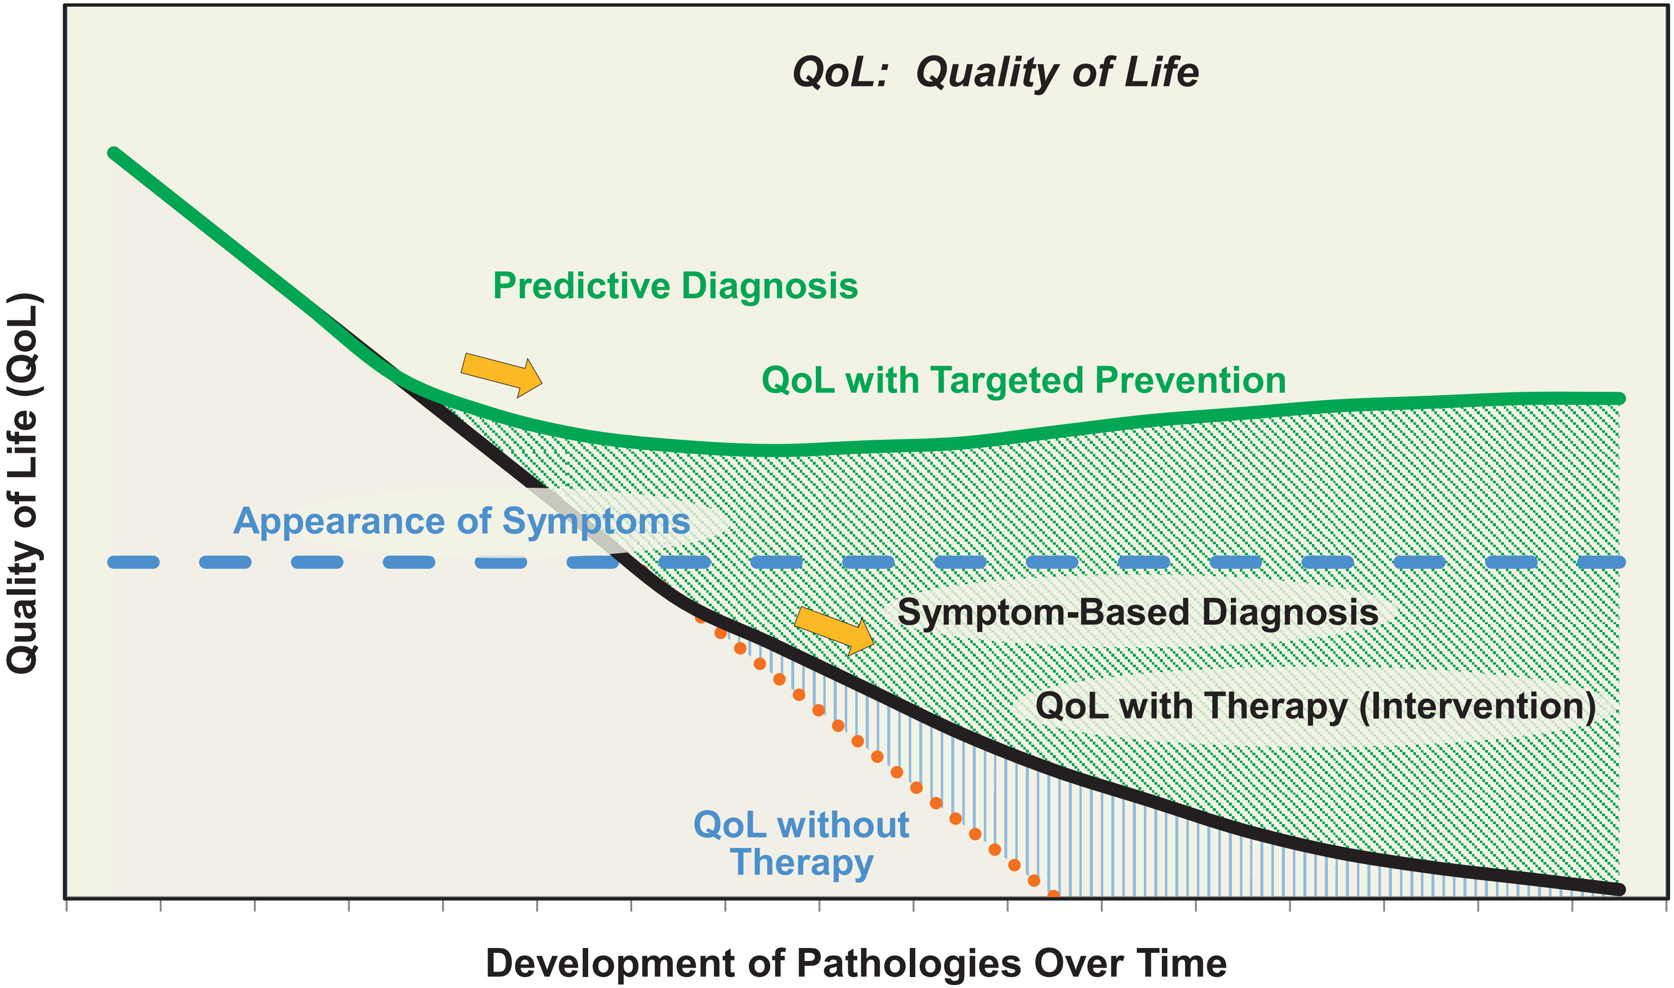

Adapted from: Costigliola V., P. Gahan, and O. Golubnitschaja: Predictive Medicine as the New Philosophy in Health Care. Predictive Diagnostics and Personalized Treatment: Dream or Reality, pp. 1-3. Edited by O. Golubnitschaja; Nova Science, 2009. (ISBN 978-1-60692-737-3)

DEC_3 / 2014 / GSD / 44687

MAYO CLINIC SPPDG

Medical Application and Low Level Algorithmic Coverage

Medical areas		*Dense linear algebra*	*Sparse linear algebra*	*Graph Algorithms*	*Frequency Analysis*	*Data Retrieval/ Filtering/ Sorting*	*Stochastic processes*	*Monte Carlo Particle methods*
Basic Bio-medical Modeling	**Disease Processes**		*X*	*X*			*X*	
	Devices / Physics	*X*			*X*		*X*	*X*
	Biology	X_{new}	X_{new}		X_{new}		*X*	*X*
Clinical science	**Population Statistics**	*Trad*	*Graph*	*Graph*	*Trad*	*Both*	*Both*	
Clinical Practice	**Image Formation**	*X*			*X*		*X*	
	Image Analytics	*X*	*X*	*X*	*X*	*X*	*X*	*X*
	Genomic Analysis	*X*				*X*		
	Decision Support		*X*	*X*		*X*	*X*	X_{new}
Health Manage-ment	**Trend Analytics**	*Trad*	*Graph*	*Graph*	*Trad*	*Both*	*Both*	
	Privacy Protection			X_{new}		*X*		

Why CnC

- Separation between
  - Domain spec (Constraint Graphs)
  - Algorithm
  - Tuning spec
- Dwarfs and Elfs map nicely
- Infinite recursion possible

(Turtles all the way down)

Of what use Dwarfs?

Of what Use Elves?

Why would anyone want to make a hybrid computing platform?

COMPARISON OF QUERY TIME FOR URIKA 2.0, MYSQL, AND SPEED_MT USING DOG_100G DATASET
(Semantic Query Times on Cray XMT Increase Linearly with Complexity; Relational Query Times are Fast Up to a Limit)

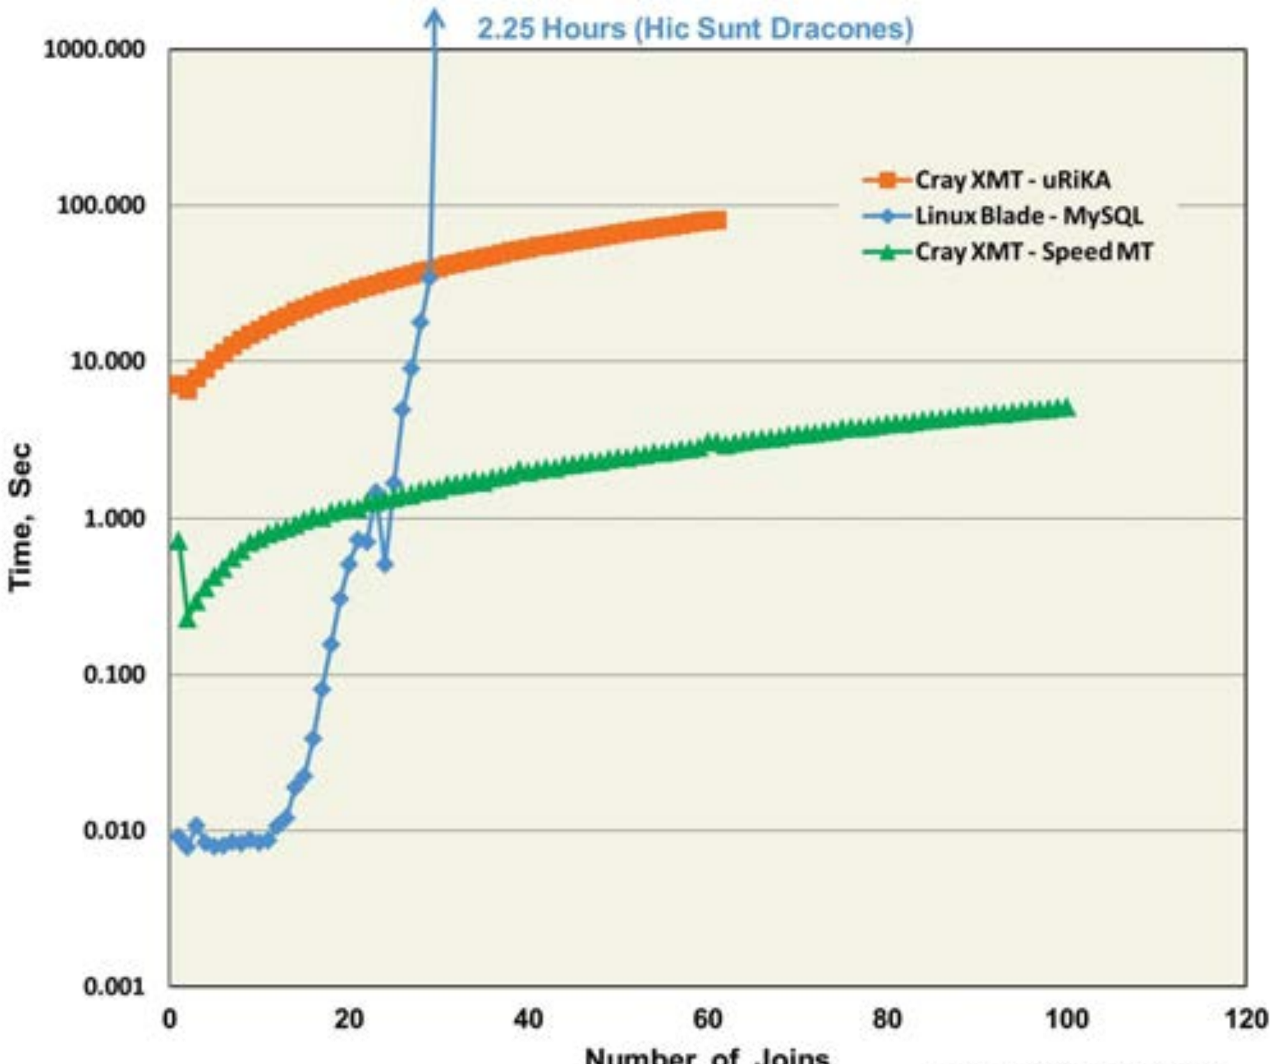

DEC_11 / 2012 / DJS / 43805
MAYO CLINIC
SPPDG

COMPUTATION RATES FOR MAXIMUM LIKELIHOOD GENOME ASSEMBLY ALGORITHM RUN ON CRAY XMT-2 AND INTEL-BASED SERVER WITH OPENMP

(Algorithm Implemented Using MTGL Hash Tables and Array Data Structures; Cray XMT-2 System with 64 Processors and 2 TB RAM; Intel Xeon Server with 4 CPUs, 32 Cores, and 0.25 TB RAM)

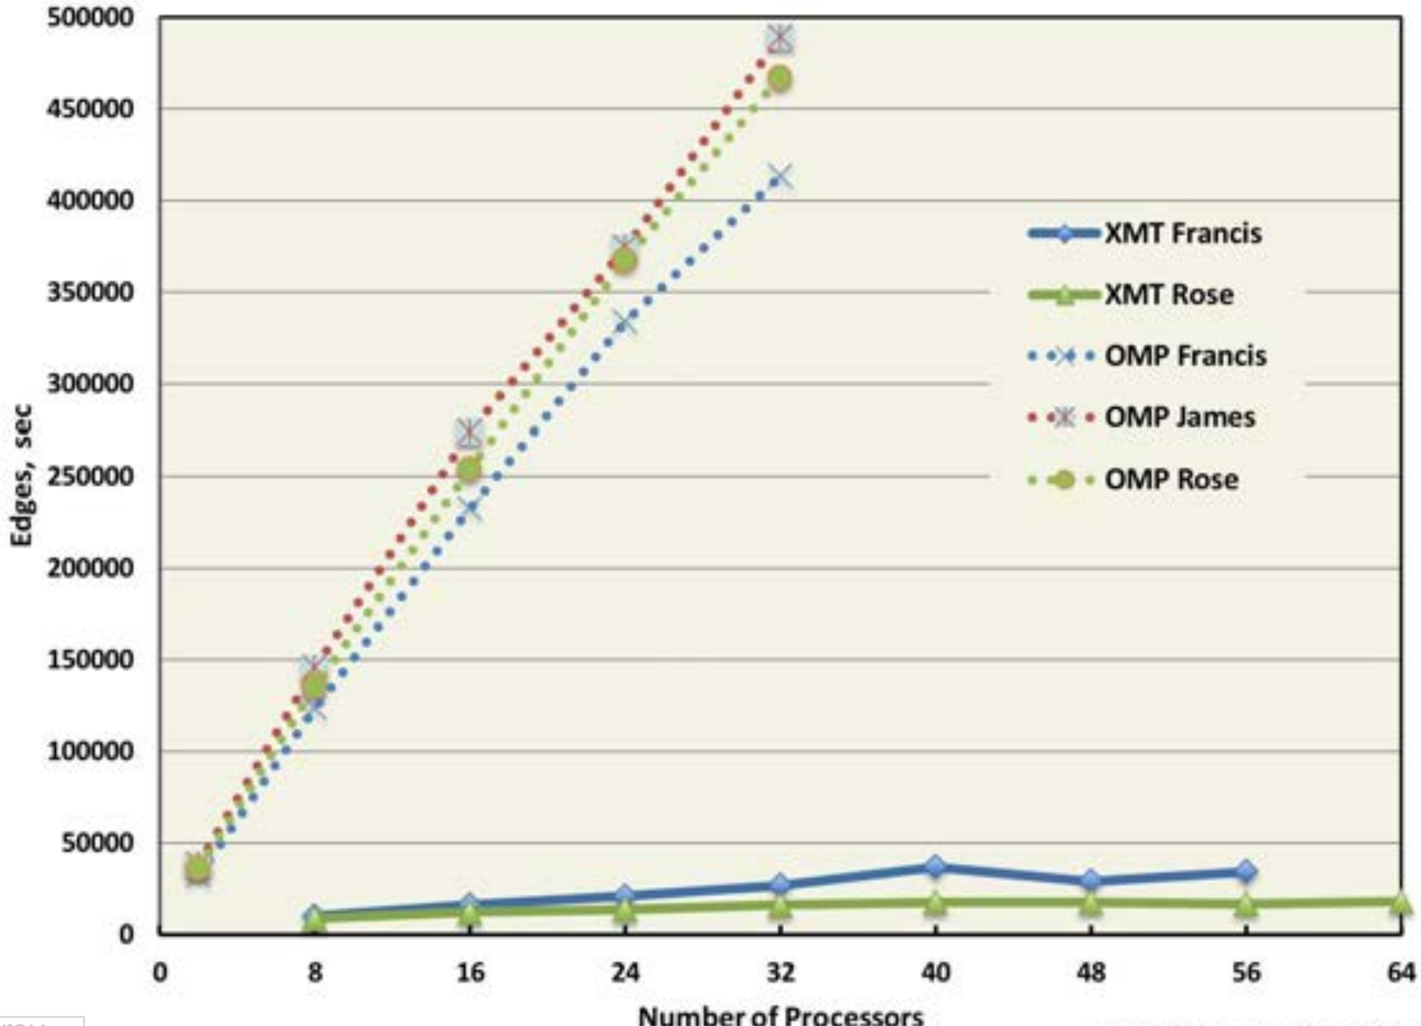

BENCHMARKS OF PLATFORMS RUNNING MOLECULAR DYNAMICS SIMULATION USING THE AMBER MD AND PREMD TOOLSETS

Simulating 23,558 Atoms comprising the Enzyme Dihydrofolate Reductase (DHFR) using Four-Femtosecond Time-Slices, in the Number, Volume, Energy (NVE) Ensemble Using Hydrogen Mass Repartitioning (HMR)

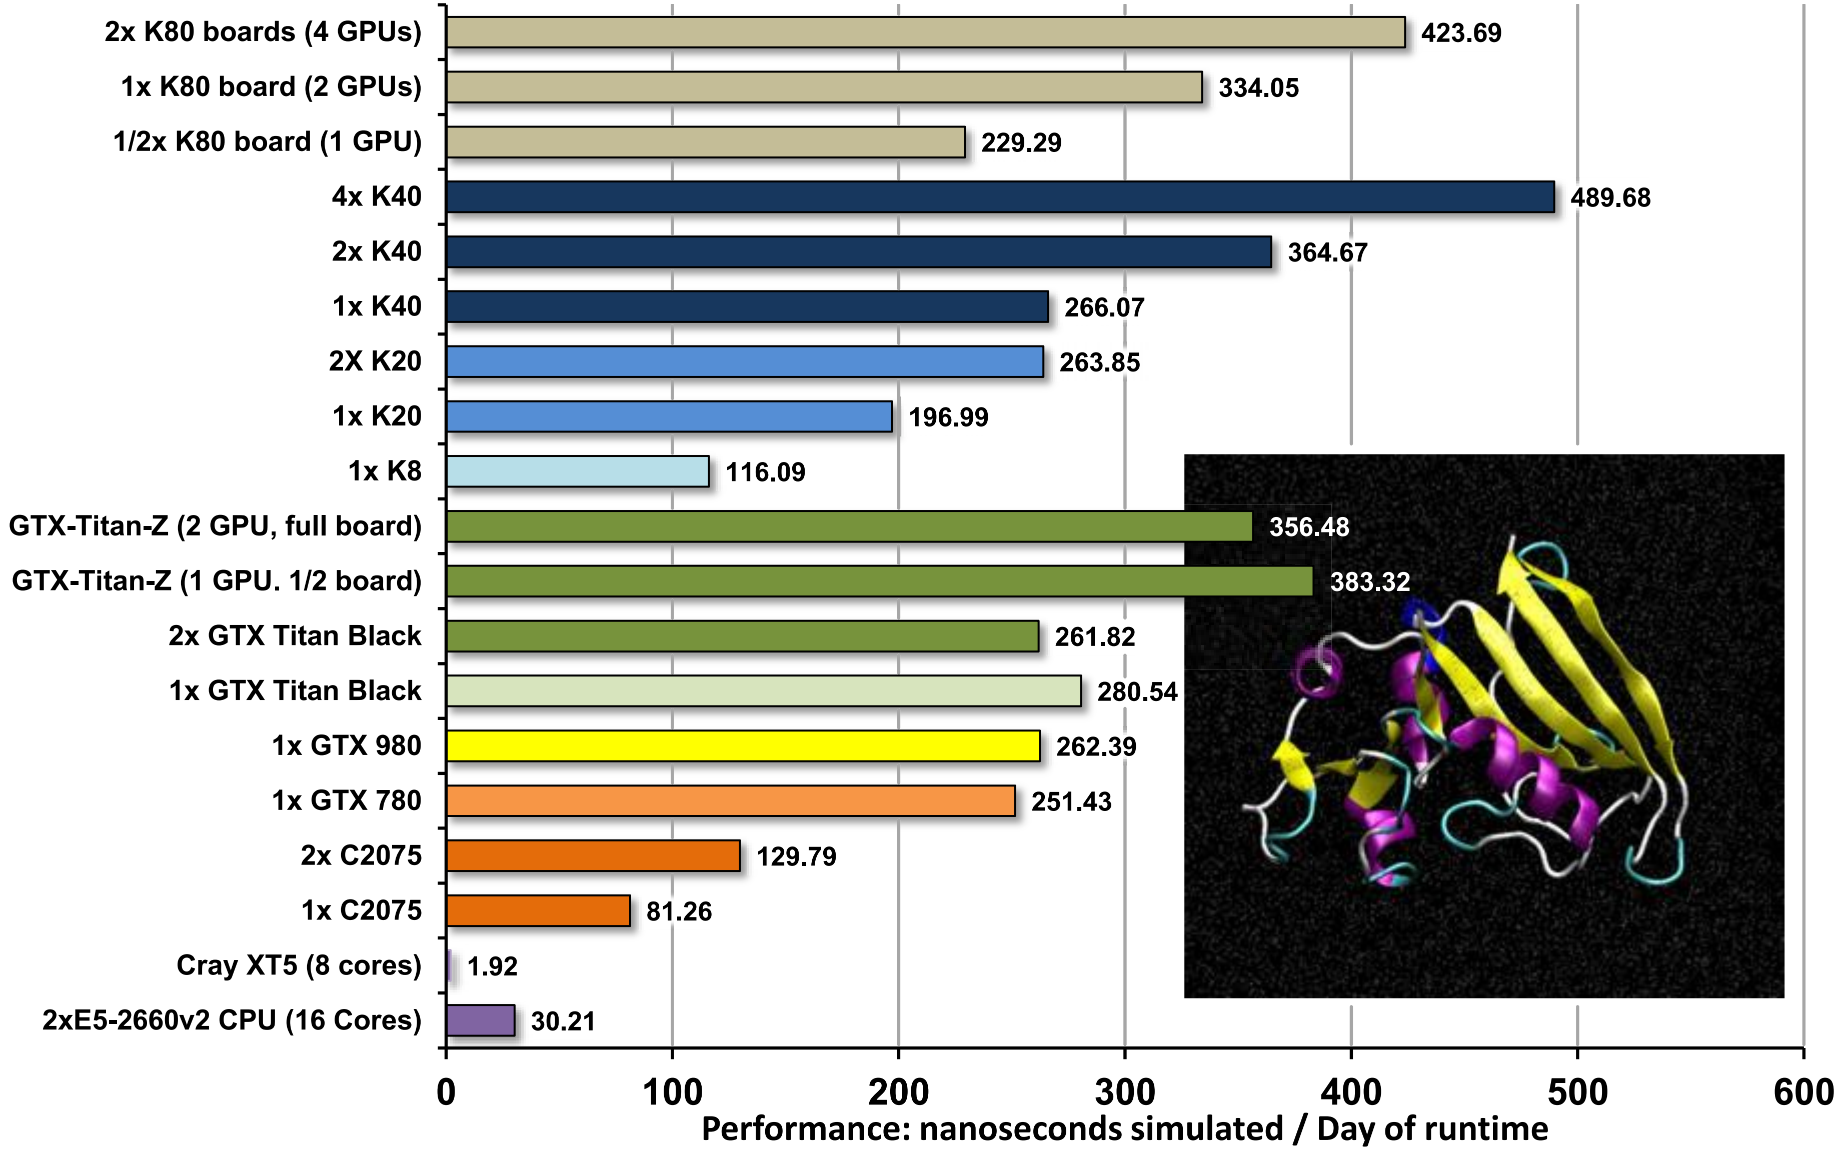

Adapted from data in, Walker, R. and S.L. Grand: *"Amber 14 NVIDIA GPU Acceleration Support"*, 2015, http://ambermd.org/gpus/benchmarks.htm & Salomon-Ferrer, R., D.A. Case and R.C. Walker: *"An Overview of the Amber Biomolecular Simulation Package,"* WIREs Comput Mol Sci, Wiley-Blackwell, 2012, *http://dx.doi.org/10.1002/wcms.1121*.

MAR_09 / 2015 / GSD / 44793r4

MAYO CLINIC SPPDG

TREATMENT PLAN EVALUATION FOR INTERSTITIAL PHOTODYNAMIC THERAPY IN A MOUSE MODEL BY MONTE-CARLO SIMULATION

(As Described by Cassidy, Betz, and Lilge; Using “FullMonte” FPGA Acceleration for the Simulation and Evaluation Steps; Each Treatment Must Be Evaluated Based On Current Parameters)

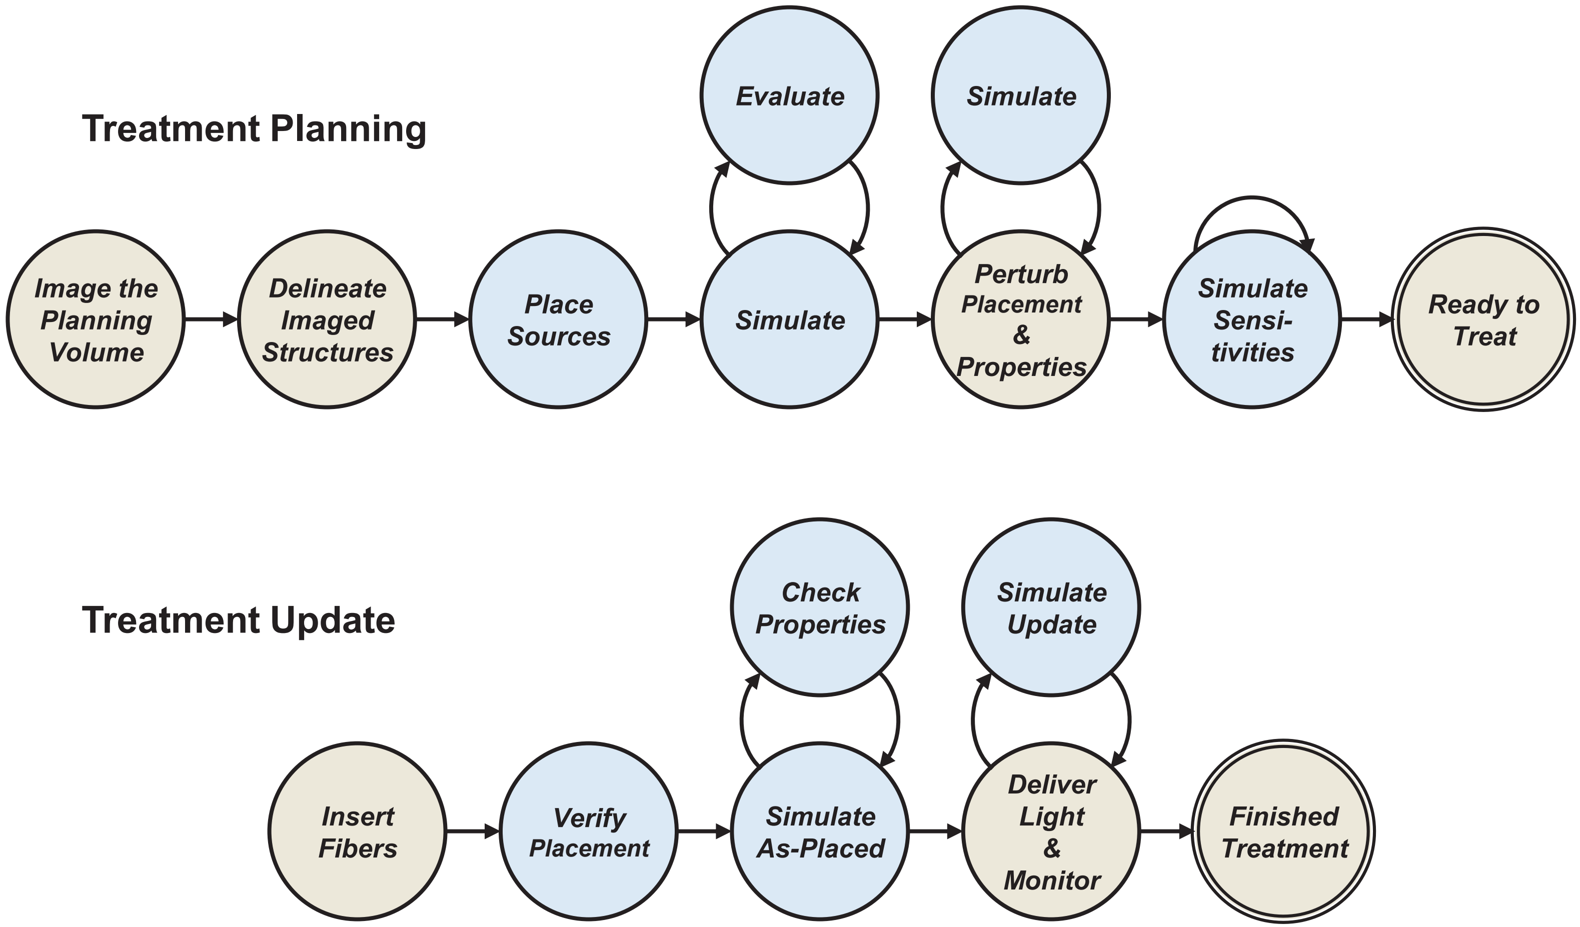

Adapted with permission from, Cassidy, J., V. Betz, and L. Lilge: “Treatment Plan Evaluation for Interstitial Photodynamic Therapy In A Mouse Model By Monte Carlo Simulation With FullMonte”. Front. Phys., 3:6 (Feb) 2015; doi:10.3389/fphy.2015.00006.

Performance and Energy-Efficiency Comparison
FPGA VS CPU

Table 27. Resources and F_{max} for Single Instance on Stratix V A7 (From [147])

Functional block	Fmax MHz	ALM	FF	DSP	M20k BRAM
Point source	290	1792	2014	2	2
Henyey-Greenstein	364	1740	2857	4	0
Scatter	302	280	546	19	0
TT800 RNG	590	804	800	0	0
Intersection test	329	510	799	20	0
Boundary	340	1707	2713	5	2
Step finish	*			3	0
Mesh storage	*			0	1034
Fluence accumulation	*			0	211
Total	280	16271	29154	59	1265
% of Available		7%	6%	23%	49%
* Not synthesized individually; no isolated Fmax available					

Table 28. Performance and Energy-Efficiency Comparison (FPGA VS CPU) (From [147])

		Relative	
Platform	Power (W)	Speed	Energy/op
CPU	76	1	67.5
Single-instance Stratix V	4.5	4	1
Estimated 4 instances	13.9	16	0.77

$$\left(1400 \approx 16 \times \frac{67.5}{0.77}\right)$$

While this problem may also experience significant speedup with the dense floating point resources of a Xeon Phi™ or a GPU, Jeffery Cassidy shared in his **SPPDG ROLEX** presentation that this is a cache unfriendly computation (Xeon thrashes), and the a GPU's local memory does not come close to scaling with its processing capability (GPU non-starter).

Given these limitations, this FPGA **14-bit fixed point** solution remains an example where the reputation of "power-hungry" that FPGA solutions have had in the past is no longer true when solving problems that suit the FPGA's capabilities.

Evaluating the power and performance, the quad FPGA solution shows a 1400 × improvement

[147] Cassidy, J., V. Betz, and L. Lilge : "Treatment Plan Evaluation for Interstitial Photodynamic Therapy In A Mouse Model By Monte Carlo Simulation With FullMonte". Front. Phys., 3:6 (Feb) 2015; doi:10.3389/fphy.2015.00006 .

PROCESSING TRADEOFFS IN A WORKFLOW VARY OVER TIME

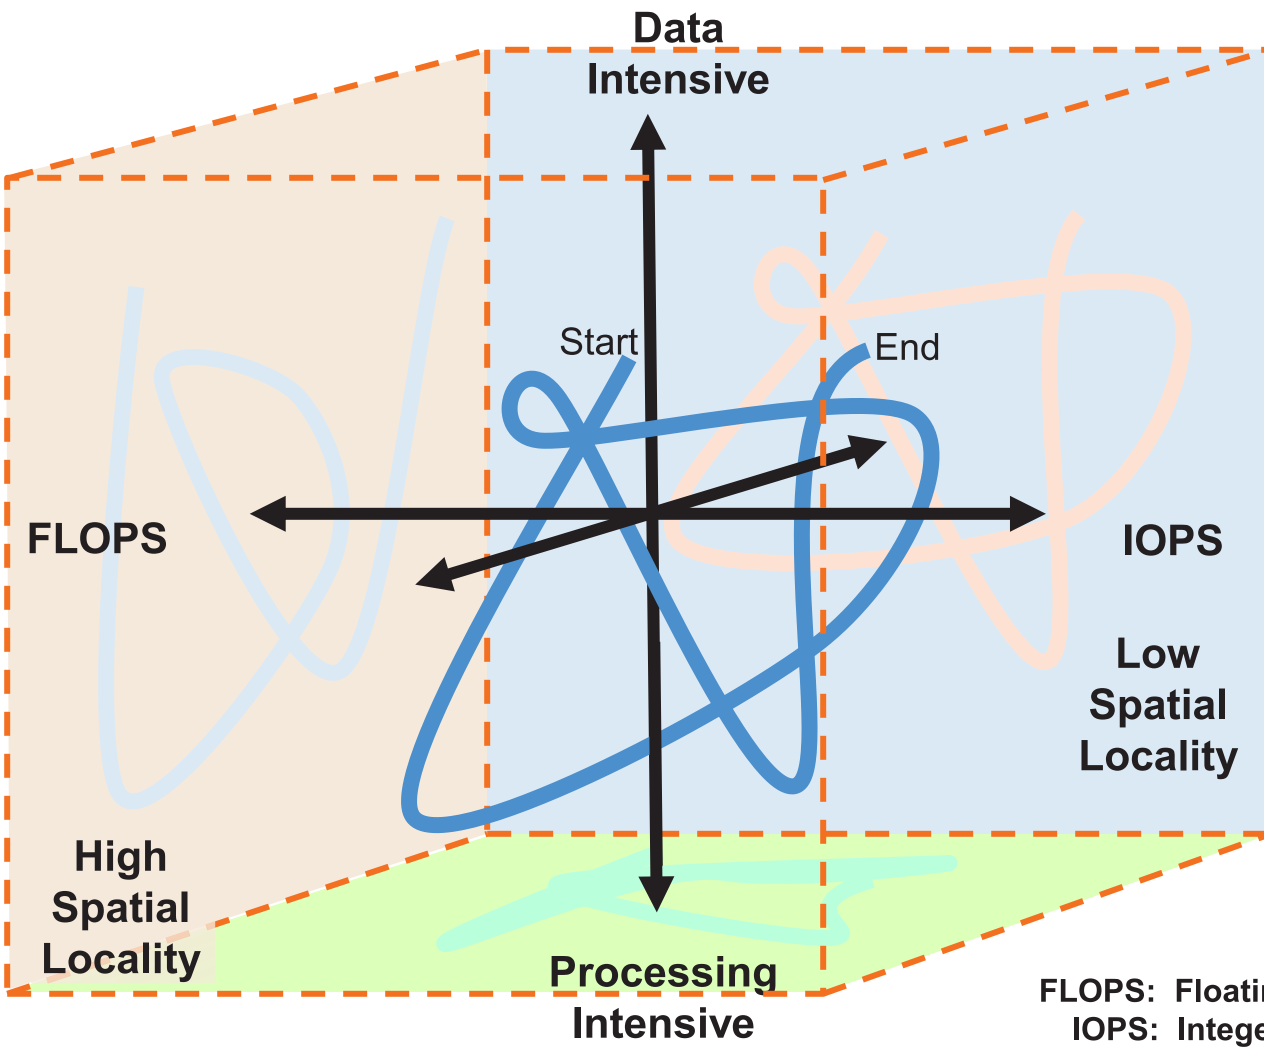

FLOPS: Floating Point Operations
IOPS: Integer Operations

Adapted with permission from: Pradip Bose, Energy Efficiency and Resilience Tradeoffs: Architecture and Modeling Challenges, Supercomputing 2013.

THE VIEWS THAT CAN BE TAKEN OF BIG DATA OGRES AND THE FACETS OF THOSE VIEWS
(The *Views* Include Data Source and Style, the Problem Architecture, Execution, and the Processing View; This Is From Early Work By Fox, et al., On Developing A Systematic Approach To *Big Data* Benchmarking)

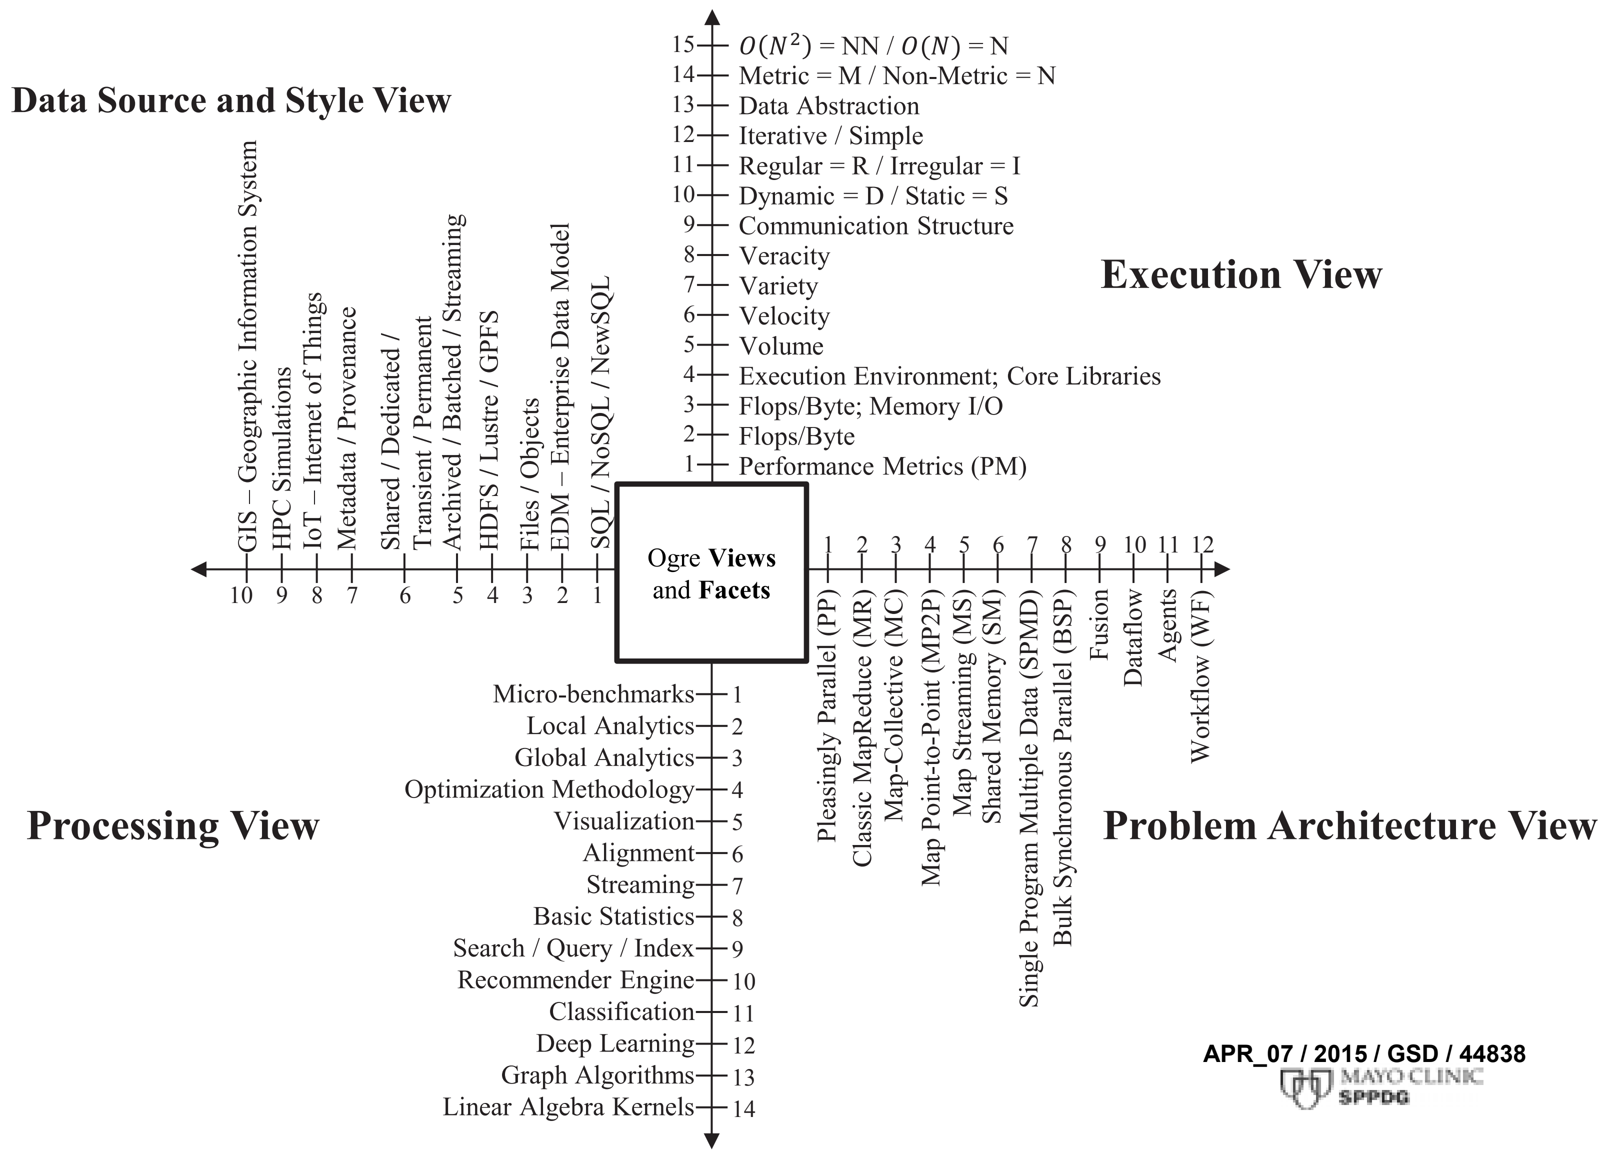

APR_07 / 2015 / GSD / 44838

MAYO CLINIC SPPDG

Adapted from, Fox, G.C., et al.: "Towards a Systematic Approach to Big Data Benchmarking," Community Grids Lab: Pervasive Technology Labs, Computer Science and Informatics, Indiana University, Bloomington, IN, Technical Report submitted for publication, 15 February 2015; ***http://grids.ucs.indiana.edu/ptliupages/publications/OgreFacetsv9.pdf***

SYNTHESIS-BASED COMPUTER-AIDED ENDOPEPTIDASE MOLECULAR DESIGN OF A POTENT NEW SMALL-MOLECULE INHIBITOR OF THE NEUROTOXIN BOTULINUM

(Illustration is the Large Enzyme-Substrate Interface of BoNTAe; A: Top View of the Active Site Showing the Substrate Binding at the Large Pocket; B: Side View of the Active Site Showing the Substrate Wrapping Around the Circumference of BoNTAe; Active-Site Residues of BoNTAe (Zn^{+2}, H223, H227, E262, F163, F194, R363, and D370) are Shown In Light Blue Sphere or Light Blue Stick Model; the SNAP-25 Substrate (146 - 204) is Shown In Red Stick Model; BoNTAe is Shown In Grey Surface Model with 15% Transparency)

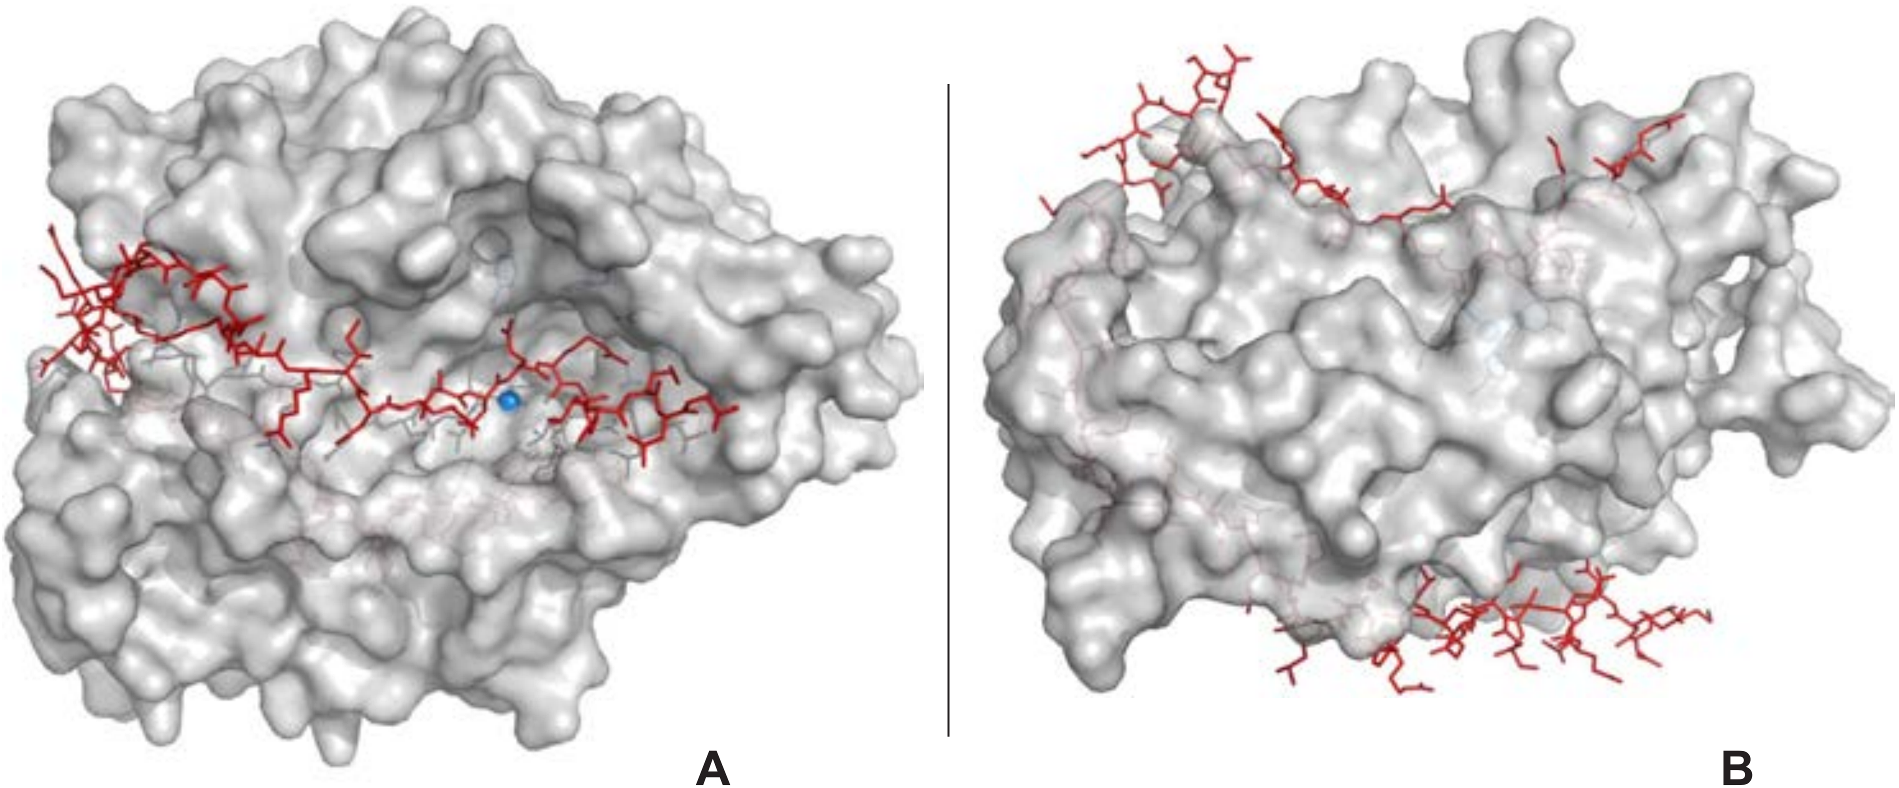

Adopted from, Pang, Y.-P., et al.: "Potent New Small-Molecule Inhibitor of Botulinum Neurotoxin Serotype a Endopeptidase Developed by Synthesis-Based Computer-Aided Molecular Design". PLoS ONE, 4(11):e7730 (2009); DOI: 10.1371/journal.pone.0007730; http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2771286/.

EXAMPLE OF DISEASE MECHANISM MODELING FROM "NOVEL AND VIABLE ACETYLCHOLINESTERASE TARGET SITE FOR DEVELOPING EFFECTIVE AND ENVIRONMENTALLY SAFE INSECTICIDES"

(Overlay of the African Malaria Mosquito (Green) and Human (Yellow) Acetylcholinesterases (Neurotransmitters) from A Perspective Looking Down Onto Substrate Acetylcholine at the Catalytic Site)

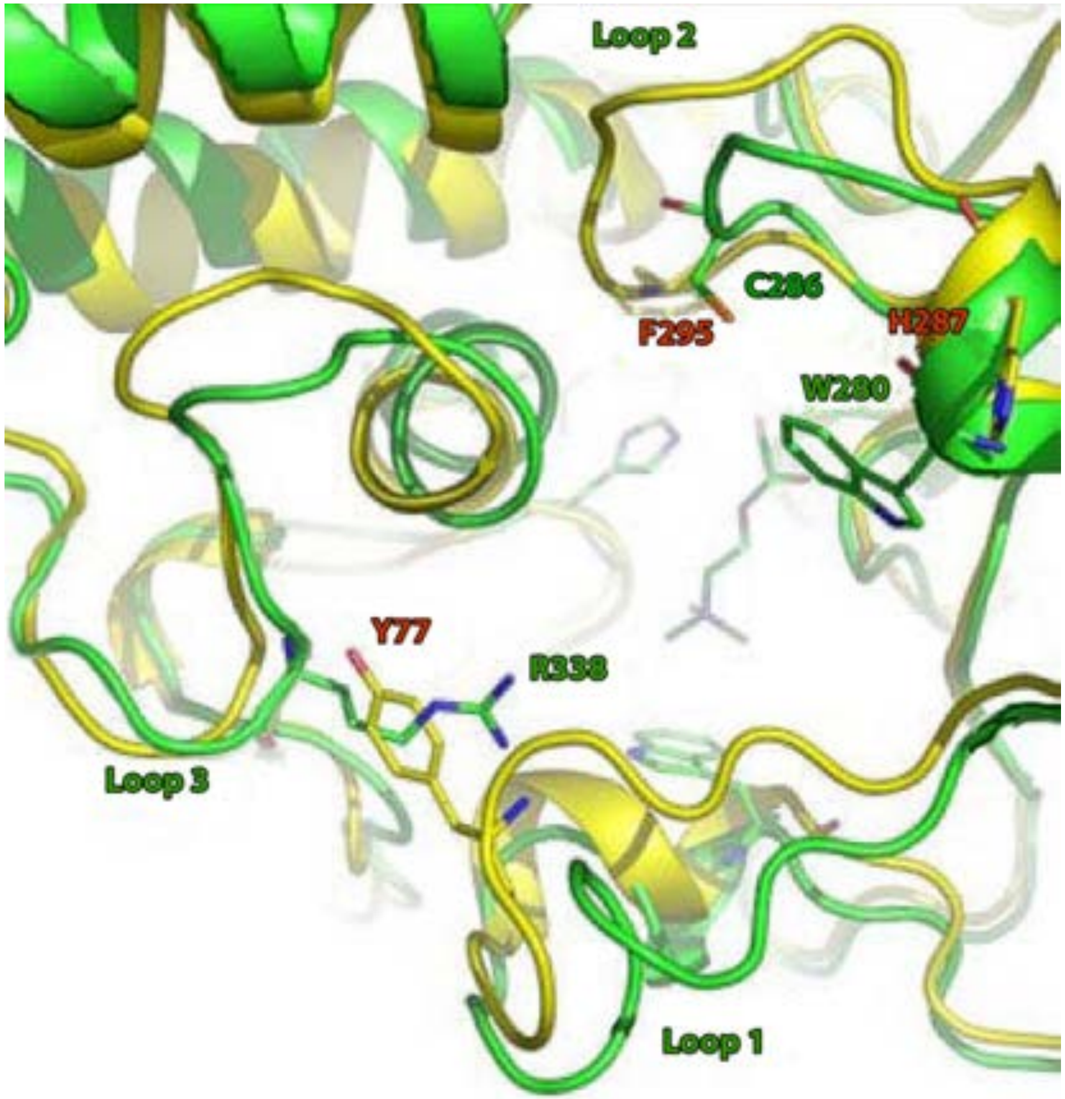

MAR_17 / 2015 / GSD / 44808

MAYO CLINIC SPPDG

Adapted from, Pang, Y., S. Brimijoin, D. Ragsdale, K. Zhu and R. Suranyi: "Novel and Viable Acetylcholinesterase Target Site for Developing Effective and Environmentally Safe Insecticides". Curr Drug Targets, 3(4):471-482 (2012); http://www.ncbi.nlm.nih.gov/pubmed/22280344.

Conclusions

- Hybrid computing is necessary to efficiently compute complex results
- Dwarfs can use CnC to manage their tile/block/search parallelism
- Elfs can use CnC to mange dwarfs
- Tuning spec: Scheduling tasks and assigning workload based not only on "the best processing element" but also on affinity of data, current workload, and other figures of merit
- The one time tagging of data may benefit from augmentation when data are produced from database queries. Interaction with adding and pruning SQL, NOSQL, and especially graph-based data stores (databases) needs exploration.
- Mayo/SPPDG streaming applications may exercise parts of the CnC abstraction that have not been implemented. Preliminary discussions reveal that there are many ways to accomplish the needs of streaming/continuous processes. Narrowing the solution space will be helpful.

Some “Wants” that may be of interest

- Transactions collections of computation steps
- Hierarchical CnC (Miland)
- Separation between algorithm and tuning
- Data dependent get
- Dynamic graph construction