Microsoft

EXTREME COMPUTING GROUP

Defining the future.

The Cloud, the Client and Big Data

Dennis Gannon
Extreme Computing Group
Microsoft Research

dennis.gannon@microsoft.com

Outline

- The Data Deluge & the 4th Paradigm
- The Cloud and Modern Data Centers
- Cloud Applications and Algorithm Challenges
- Changing the Paradigm & Democratizing Science
- The Cloud Research Engagement Project

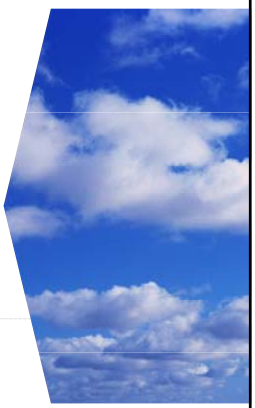

Microsoft

The Microsoft Cloud is Built on Data Centers

~100 Globally Distributed Data Centers

Range in size from "edge" facilities to megascale (100K to 1M servers)

The Microsoft Cloud

- Purpose-built data centers to host containers at large scale
  - Cost $500 million, 100,000 square foot facility (10 football fields)
- 40 foot shipping containers can house as many as 2,500 servers
  - Density of 10 times amount of compute in equivalent space in traditional data center
- Deliver an average PUE of 1.22
  - Power Usage Effectiveness benchmark from The Green Grid™ consortium on energy efficiency

Data Center Infrastructure

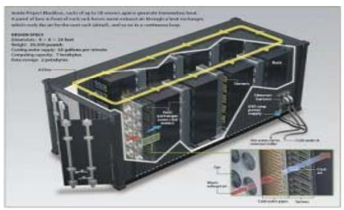

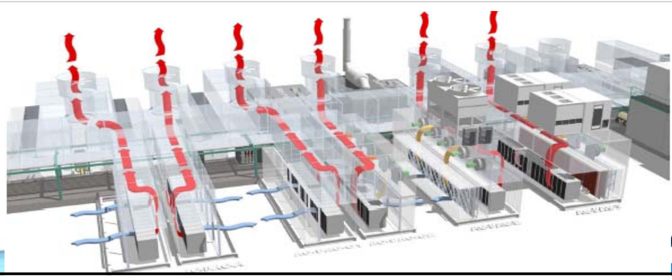

Data Center vs Supercomputers

- Scale
  - Blue Waters = 40K 8-core "servers"
  - Road Runner = 13K cell + 6K AMD servers
  - MS Chicago Data Center = 50 containers = 100K 8-core servers.
- Network Architecture
  - Supercomputers: CLOS "Fat Tree" infiniband
  - Low latency – high bandwidth protocols
  - Data Center: IP based
    - Optimized for Internet Access
- Data Storage
  - Supers: separate data farm
    - GPFS or other parallel file system
  - DCs: use disk on node + memcache + databases

Fat tree network

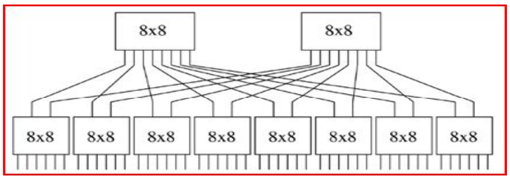

Standard Data Center Network

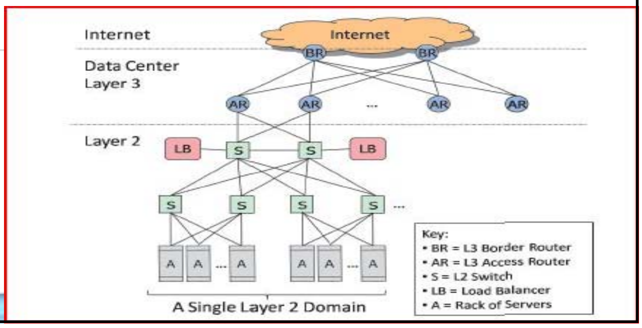

The Cloud Everybody Uses Today

- A revolution in mobile applications
  - Derive their power from cloud services
  - Maps and location services
  - Synchronized calendar and mail
  - Cloud based media storage
- A revolution in smart Internet search
  - Driving advances in big data analytics
  - A new frontier for machine learning
    - Speech and language translation

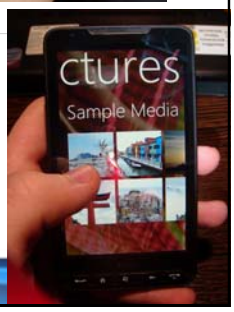

The 4th Paradigm and the Revolution in Science

The Data Explosion

Experiments | Simulations | Archives | Literature | Consumer

The Challenge
Enable Discovery

Deliver the capability to mine, search and analyze this data in near real time

Petabytes
Doubling & Doubling

The Response

Science itself is evolving

The Changing Nature Of Research

$$\left(\frac{a}{a}\right)^2 = \frac{4\pi G\rho}{3} \quad K\frac{c^2}{a^2}$$

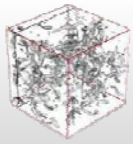

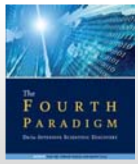

Experimental	Theoretical	Computational	The Fourth Paradigm
Thousand years ago *Description of natural phenomena*	Last few hundred years *Newton's laws, Maxwell's equations…*	Last few decades *Simulation of complex phenomena*	Today and the Future *Unify theory, experiment and simulation with large multidisciplinary Data* *Using data exploration and data mining (from instruments, sensors, humans…)* *Distributed Communities*

Microsoft

Some Examples of 4th Paradigm Applications

- Genomics
  - Build the phylogenic tree that describes the evolution of AIDS
- Text Analysis
  - Natural language translation
  - Mine and map the key concepts in a large corpus of text
- Social graph analysis
- Web image clustering and analysis
- Computer Vision (Kinect)

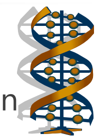

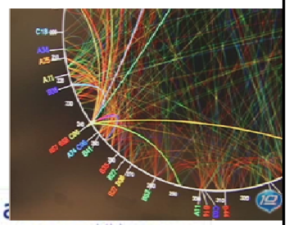

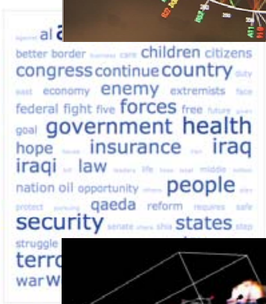

What Do These Applications Have in Common?

- **A massive data analytics phase based on parallel algorithms working on distributed data.**
  - Data is partitioned and stored across a very large number of servers
    - Often replicated. This is a NoSQL database
  - The algorithms are often based on a MapReduce parallelism
- **Web services that allow user to use the cloud as an extension of the desktop/client device**
  - Manage the distributed data collections
  - Invoke data analytics algorithms
  - Visualize the results
  - Share and collaborate

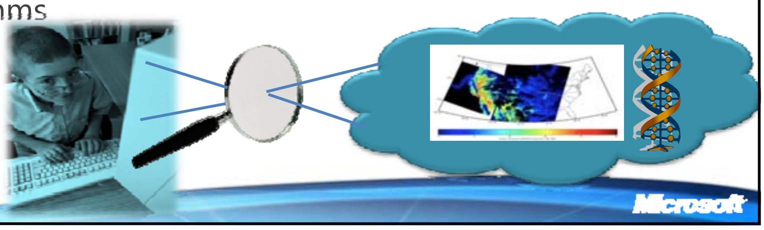

Cloud Properties

- **Designed to Provide Information and Computation to Many Users**
- Automatic Deployment and Management of Virtual Machine Instances
  - tens to thousands & dynamic scalability
- Dynamic Fault Recovery of Failed Resources
  - Cloud Services must run 24x7
- Automatic Data Replication
  - Geo-replication if needed
- Two levels of parallelism
  - Thousands of concurrent users
  - Thousands of servers for a single task.

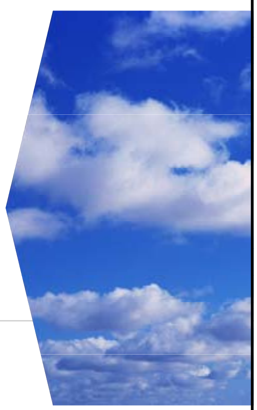

The Nature of Cloud Data Applications

- The rise of MapReduce
  - How do you build an Index of the web?
  - How to you find clusters in a billion images?
  - How do you invent a Bayesian net for predicting customer behavior from the twitter fire hose?
  - Or phylogenic/evolutionary relationships in gene sequences
- noSQL, distributed column stores superior for petascale data analytics

Cloud Computing Requires a Different Model

- Traditional Model of Computing
  - You load an application into your operating system on your machine and run it.
- The Cloud Model
  - Your browser or desktop application uses remote data and services to do your computation
- As cloud programmers you should
  - Think about building application services that run forever and solve many instances of a problem
  - Data is massive and distributed across thousands of disks or memories
  - Failure is the norm. Computation should be stateless and redundant

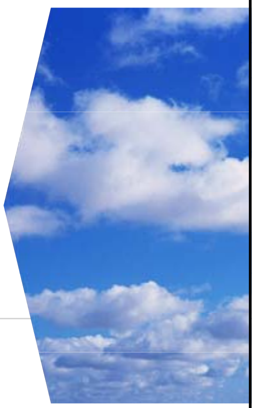

Cloud Algorithm Challenges

- Parallel algorithms for massively distribute data
  - ❑ Clustering: K-means, fuzzy clustering, canopy clustering;
  - ❑ Recommendation Mining: Log-Likelihood;
  - ❑ Prediction: SVM; trend prediction
  - ❑ Frequent Item Set Mining: Collocation, Outlier Detection;
  - ❑ Bayesian/Regression Toolkit (linear, non-linear, logistic);
  - ❑ Bayesian Net, Neural Nets, other Machine learning
- Grahps, Tables, Matricies
  - ❑ Eigenvalues/eigenvectors are the key:
    - PageRank is a simple Principal Eigenvector problem
    - Latent Semantic Indexing (LSI), Principal Component Analysis (PCA) is another example
    - Clustering is accomplished by taking the top k eigenvectors
    - Partitioning is done by taking the bottom k eigenvectors
    - Feature extraction/comparison can be done by SVD
  - ❑ Graph partitioning / spectral clustering critical for storage optimization

Some Examples of 4th Paradigm Applications

- Genomics
  - Build the phylogenic tree that describes the evolution of AIDS
- Text Analysis
  - Natural language translation
  - Mine and map the key concepts in a large corpus of text
- Social graph analysis
- Web image clustering and analysis
- Computer Vision (Kinect)

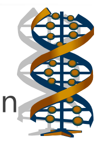

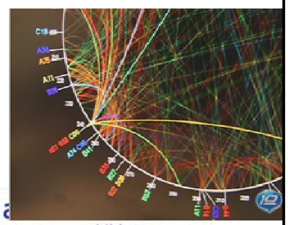

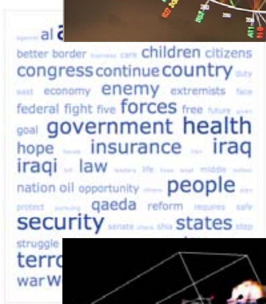

What Do These Applications Have in Common?

- **A massive data analytics phase based on parallel algorithms working on distributed data.**
  - Data is partitioned and stored across a very large number of servers
    - Often replicated. This is a NoSQL database
  - The algorithms are often based on a MapReduce parallelism
- **Web services that allow user to use the cloud as an extension of the desktop/client device**
  - Manage the distributed data collections
  - Invoke data analytics algorithms
  - Visualize the results
  - Share and collaborate

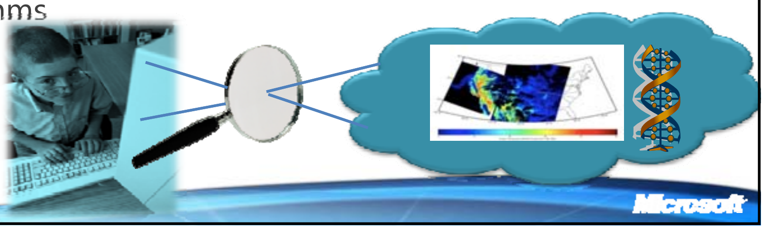

Programming the Cloud: Cloud Properties

- **Data Centers Designed to Provide Information and Computation to Many Users**
- Automatic Deployment and Management of Virtual Machine Instances
  - tens to thousands & dynamic scalability
- Dynamic Fault Recovery of Failed Resources
  - Cloud Services must run 24x7
- Automatic Data Replication
  - Geo-replication if needed
- Two levels of parallelism
  - Thousands of concurrent users
  - Thousands of servers for a single task.

Cloud Computing Requires a Different Model

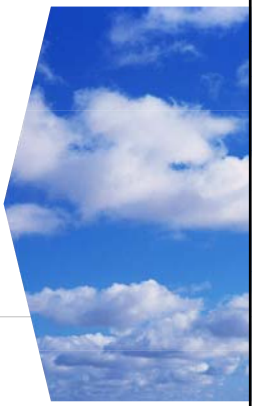

- Traditional Model of Computing
  - You load an application into your operating system on your machine and run it.
- The Cloud Model
  - Your browser or desktop application uses remote data and services to do your computation
- As cloud programmers you should
  - Think about building application services that run forever and solve many instances of a problem
  - Data is massive and distributed across thousands of disks or memories
  - Failure is the norm. Computation should be stateless and redundant

The Nature of Cloud Data Applications

- The rise of MapReduce
  - How do you build an Index of the web?
  - How to you find clusters in a billion images?
  - How do you invent a Bayesian net for predicting customer behavior from the twitter fire hose?
  - Or phylogenic/evolutionary relationships in gene sequences
- noSQL, distributed column stores superior for petascale data analytics

MapReduce

- Parallel data analytics algorithms such as clustering, machine learning and more. For example build an index of the Web.
- Step 1. upload or extract the data set as a distributed collection
- Step 2. apply a user provided function F to each data item
- Step 3. reduce with user provided associative operator o to reduce the set of results from 2 to a single summary

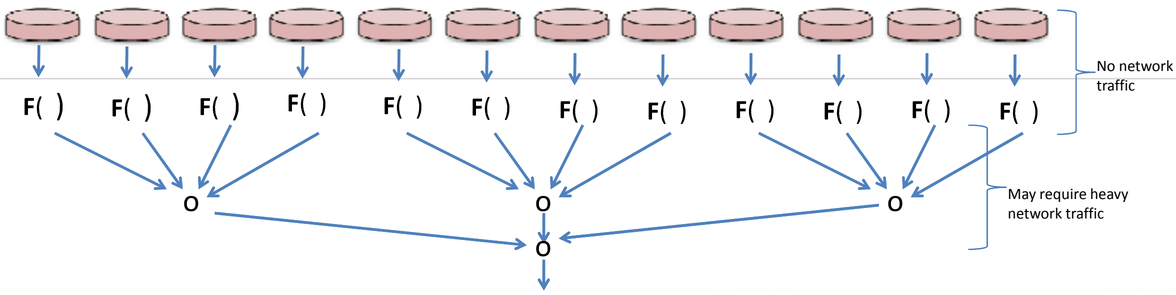

Cloud Algorithm Challenges

- Parallel algorithms for massively distribute data
  - ❑ Clustering: K-means, fuzzy clustering, canopy clustering;
  - ❑ Recommendation Mining: Log-Likelihood;
  - ❑ Prediction: SVM; trend prediction
  - ❑ Frequent Item Set Mining: Collocation, Outlier Detection;
  - ❑ Bayesian/Regression Toolkit (linear, non-linear, logistic);
  - ❑ Bayesian Net, Neural Nets, other Machine learning

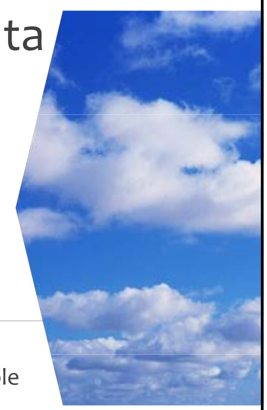

- Graphs, Tables, Matrices
  - ❑ Eigenvalues/eigenvectors are the key:
    - PageRank is a simple Principal Eigenvector problem
    - Latent Semantic Indexing (LSI), Principal Component Analysis (PCA) is another example
    - Clustering is accomplished by taking the top k eigenvectors
    - Partitioning is done by taking the bottom k eigenvectors
    - Feature extraction/comparison can be done by SVD
  - ❑ Graph partitioning / spectral clustering critical for storage optimization

Some examples based on our ongoing Cloud Research Engagements

NCBI Blast on Windows Azure

Research

- National Center for Biotechnology Information algorithm for comparing biological sequences
- Using the cloud, University of Washington reduced compute time on a project on protein interactions from 6 years to 1 week.
- Accelerating scientific discovery

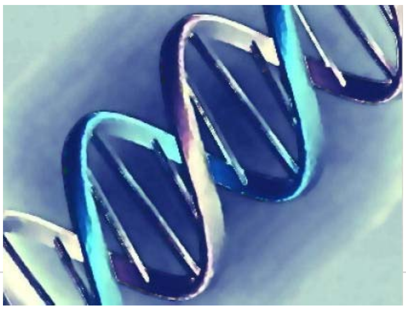

NCBI Blast for Azure

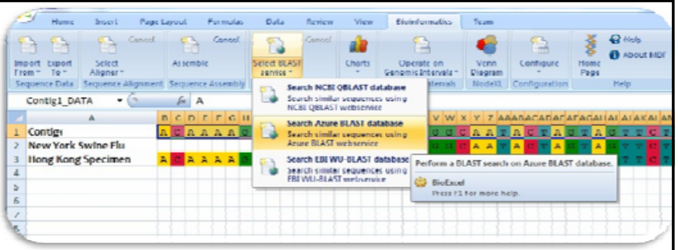

Seamless Experience

- Evaluate data and invoke computational models from Excel.
- Computationally heavy analysis done close to large database of curated data.
- Scalable for large, surge computationally heavy analysis.
- Test local, run on the cloud.

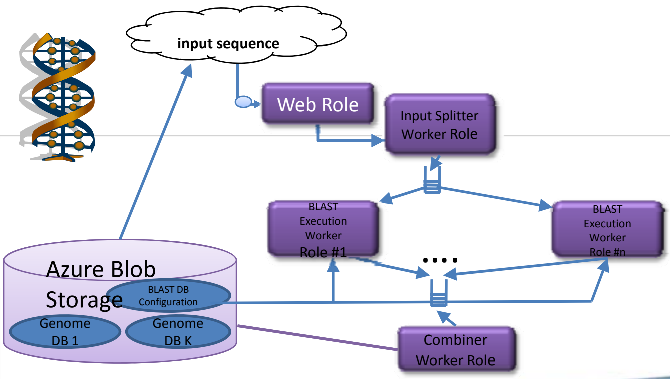

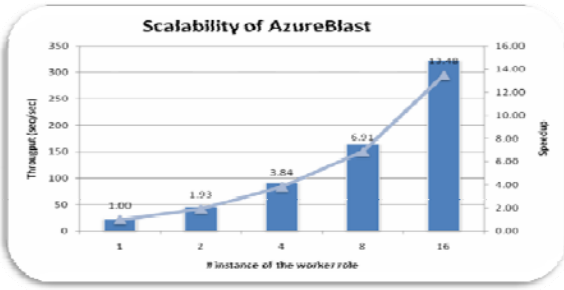

Now available for public release
http://research.microsoft.com/azure

Making Excel the user interface to the cloud

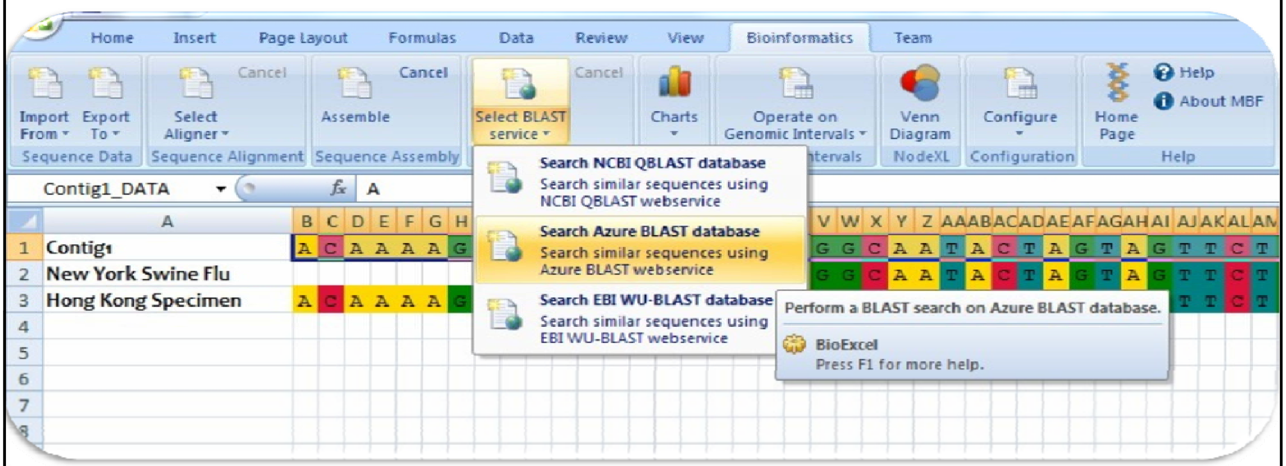

Supporting Smart Sensors and Data Fusion

- The NSF Ocean Observing Initiative
  - Hundreds of cabled sensors and robots exploring the sea floor
  - Data to be collected, curated, mined
  - OOI Architecture plan of record, store this data in the cloud

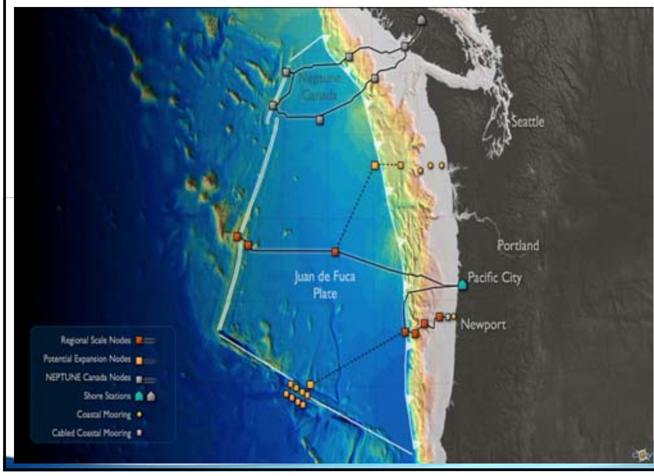

Data collected from:
- Ocean floor sensors, AUV tracks, ship-side cruises, computational models

Data moves from **ocean** to shore side **data center** to the **Azure cloud** to your **computer.**

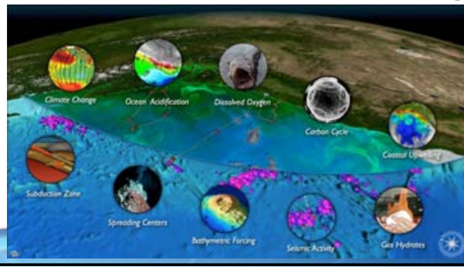

AzureMODIS – *Azure Service for Remote Sensing Geoscience*

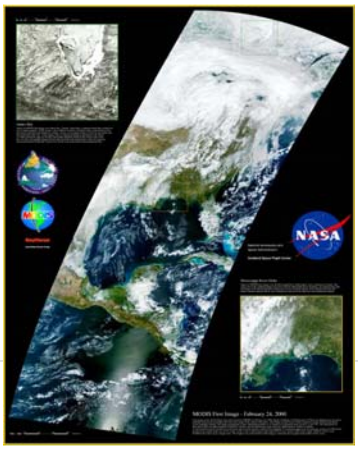

Catharine van Ingen
And her friends.

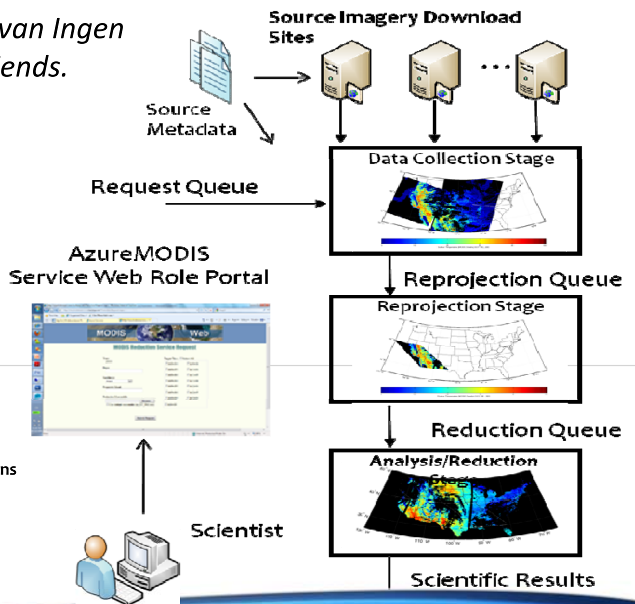

5 TB (~600K files) upload of 9 different imagery products from 15 different locations (~6 days of download)

4 TB reprojected harmonized imagery ~35000 cpu hours

50 GB reduced science variable results ~18000 cpu hours (~14 hour download)

50 GB additional reduced science analysis results ~18000 cpu hours (~14 hour download)

OPeNDAP Server on Azure

- Commonly used across the sciences
  Oceanography, Atmospheric, Earth
- Model data and sensor data
- Large communities and research projects
  OOI, JPL, PRISM, MBARI,…

Working with the OpENDAP group on the implementation

- Targeting new features for the server;
- Leverage cloud storage infrastructure for higher reliability/performance
- Identified data collections and research projects to evaluate this implementation, when it is available.

Excel DataScope

Cloud Scale Data Analytics from Excel

Bringing the power of the cloud to the laptop, familiar yet differentiated…

- **Data sharing in the cloud**, with annotations to facilitate discovery and reuse;
- **Sample and manipulate** extremely large data collections in the cloud;
- **Top 25 data analytics algorithms**, through Excel ribbon running on Azure;
- **Invoke models**, perform analytics and visualization to gain insight from data;
- **Machine learning** over large data sets to discover correlations;
- **Publish** data collections and visualizations to the cloud to share insights;

Researchers use familiar tools, **familiar but differentiated**.

How the can we make the resources and the potential for scientific discovery created by the 4th paradigm available to more people?

The Cloud – an Opportunity

Researchers are using desktop, available clusters and HPC providers

The Cloud can provide the scale, access to data and a platform for collaboration

Multidisciplinary collaboration is becoming essential

Simple Tools To Answer Complex Questions…

Imagine: the client plus the cloud for extreme data analysis

Give the standard science desktop tools a seamless extension to the cloud.

Use a spreadsheet to invoke genomic analysis tools running on 600 servers

Use a simple script to orchestrate data analytics and mining across 10000 MRI Images

Pull data from remote instruments to the cloud for visualization on the desktop

Create a revolution in scientific capability for everybody.

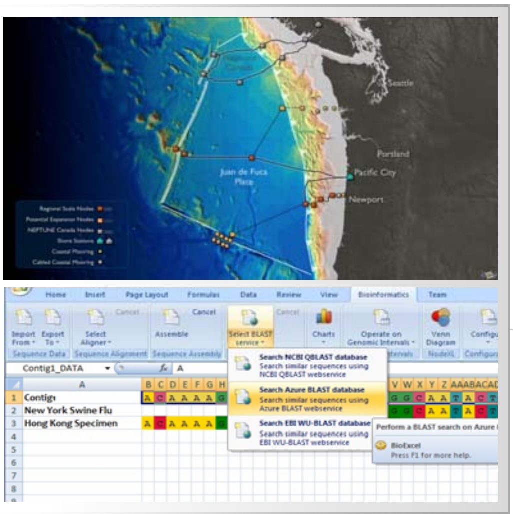

Microsoft

Let Scientists Be Scientists…

Most scientists do not want to be system administrators
They don't want to learn to use supercomputers

They want to focus on their science

They use standard tools: spreadsheets, statistical packages, desktop visualization
Programming = modifying a few parameters in a trusted scripting language

They want to share experiments with their collaborators

Microsoft

Microsoft's Global Cloud Research Engagement Initiative

Mission

- Broaden research capabilities, foster collaborative research communities and accelerate scientific discovery globally
- Provide massively scalable tools and services directly to users (from their desktops) which could transform how research is conducted, accelerating scientific exploration, discovery and results.

Approach

- Build partnerships with government-sponsored research agencies and university consortia
- Offer cloud services to academic and research communities worldwide supported with a technical engagement team.
- Provide code samples and templates for significant cloud application
- Provide tools to access the cloud and important cloud services from the desktop and laptop

Engagement Projects So Far (75)

Press Release 11-082

The Sky Is No Limit: 13 Research Teams Compute in the Clouds

NSF awards nearly $4.5 million to innovative projects to participate in NSF / Microsoft cloud computing collaboration

Cloud computing expands scientific exploration and discovery.

- **NSF Computing in the Cloud**
  - 30 projects. Areas include genomics, phylogeny, algorithms for big data, sociology, semantic web, molecular dynamics, string theory, geo ….
- **Japan National Institute for Informatics**
  - InfoPlosion : 7 projects around natural language and web analytics
- **Australia – NICTA, National Comp Infrastructure, CSIRO**
  - 6 projects: Social network analysis, realtime data, computational chemistry, magnetotellurics, coal industry logistics
- Future – More Asia Engagements
- Additional – Massive scale protein folding, genomics, n-gram services, image analysis,

Venus-C & Microsoft Demonstrating Interoperability

VENUS-C (**V**irtual multidisciplinary **E**nviro**N**ments **US**ing **C**loud infrastructures) is funded by the European Commission as part of the framework programme seven to demonstrate Cloud interoperability

Programme partners
Engineering Ingegneria Informatica SpA , Barcelona Supercomputing Centre, Centre for Computational and Systems Biology, Collaboratorio, Microsoft Innovation Center Greece, Microsoft Research, National Research Council of Italy , Royal Institute of Technology, Technion, The European Chapter of the Open Grid Forum, The European Microsoft Innovation Center, Universidad Politechnica de Valencia, University of Newcastle, University of the Aegean

Interoperability will be demonstrated around several key application areas
Biomedicine (Bioinformatics, System Biology and Drug Discovery), Civil Engineering, Data for Science, Civil Protection and Emergencies

Venus-C

Microsoft

Open Call

15 new funded partners and 11 additional resource experiments

Microsoft

Sample Projects

- **Protein Folding**
  - Post-doctoral researcher Nikolas Sgourakis of Baker Lab at Univ. of Washington is studying ways in which proteins from Salmonella fold.
  - Used 2000 concurrent cores on Azure
- **Civil Protection & Emergencies**
  - This project address is the fire risk probability prediction challenge which takes as input meteo and geo-data and calculates the so-called fire risk index.
  - The application scenario involves a client application for the fire and forest services as well as cloud services that allow access to real-time data from sensors and on-the-ground reporting.

Protein Folding

The University of Washington Baker Laboratory

Civil Protection & Emergencies

University of the Aegean

Microsoft

More Samples

- **Magnetotellurics**
  - Craig Mudge and Graham Heinson of the University of Adelaide have co-designed a platform technology for the geophysical imaging method of magnetotellurics (MT) with processing in the cloud.
  - Their MT platform applies to multiple needs: geothermal exploration, minerals exploration, water management in mining, aquifer management for indigenous communities, and carbon sequestration.
- **fMRI Brain Image Analysis**
  - .Joint genetic and neuroimaging data analysis on large cohorts of subjects is a new approach used to assess and understand the variability that exists between individuals. Both neuroimaging- and genetic-domain observations represent a large number of variables (of the order of 10^6), performing statistically rigorous analyses on such large quantities of data represents a computational challenge that cannot be addressed with conventional computational techniques.
- **Language Analysis**
  - . This team applied a predicate-argument structure analysis to a huge Japanese corpora consisting of about 20 billion web sentences, to improve the open-search engine infrastructure TSUBAKI, which is based on deep natural language processing. To achieve this goal 10,000,000 core on Windows Azure were used in a massively parallel computation that took about a week.

Development of a Magnetotellurics Platform

Using Azure at the University of Adelaide

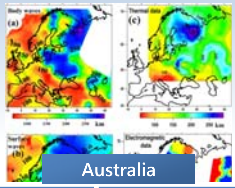

Australia

INRIA "Azure Brain"

Premier information technology research laboratory

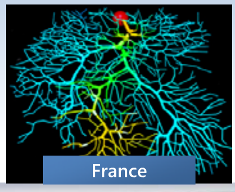

France

Infoplosion

Predicate-Argument Structure Analysis at Kyoto University

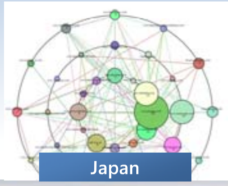

Japan

Conclusions

- There is a great potential for the Cloud, when coupled with smart desktop and mobile clients to change the way we do research.
- But there is much to learn.
  - What are the scientific applications and data collections that work best with the cloud?
  - What modes of interaction & visualization of big data work best?
  - How can research communities learn to use the cloud for collaboration?

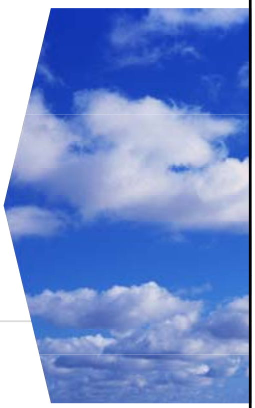

Q&A

Sample Engagement Projects

- **Inferring Pattern and Processes of Genome Evolution through Cloud Computing.** John Burleigh, University of Florida.
- **GIS Vector Data Overlay Processing on Azure Platform.** Sushil Prasad, Georgia State University
- **A String Cartography.** Brent Nelson, Northeastern University.
- **Porting the Structure-Adaptive Materials Prediction Project to the Microsoft Azure Platform.** Iosif Vaisam, George Mason University.
- **Cooperative Developer Testing with Test Intentions.** Tao Xie, North Carolina State University.
- **Towards automated and assurable enterprise network migration**. Sanjay Rao, Purdue University.
- **Data Intensive Grid Computing on Active Storage Clusters**. Douglas Thain, University of Notre Dame.
- **Moving Polarizable Force Field Simulations to the Microsoft Azure Platform**. Sandip Patel, University of Delaware.
- **Maximizing the Utility of Orthologs and Phylogenetic Profiles for Systems-Scale Comparative Genomics**. Dennis Wall, Harvard
- **Web-scale Language Modeling Features for Machine Translation**. Alon Lavie, CMU.
- **Stork Data Scheduler for Azure.** Tevfik Kosar, Louisiana State University.
- **Exploring Social Classification on Microsoft Azure.** Harris Wu, Old Dominion University.
- **Where the Ocean Meets the Cloud: Ad Hoc Longitudinal Analysis and Collaboration Over Massive Mesh Data.** Bill Howe, University of Washington.
- **Transforming Morphological Ssystematics From Desktop to Web Applications: Development of the Online Workspace Morphobank.org 3.0.** Maureen O'Leary, StonyBrook

- **Semantic Web Informatics for Species in Space and Time**. Jennifer Dunne. Santafe Institute.
- **Building Scalable Trust in Cloud Computing**. Kenneth Birman, Cornell University.
- **Bettering Interactive Protein-Protein Docking**. Andrey Tovchigrechko, J. Craig Venter Institute, Inc.
- **Enhancing Stork Data Scheduler for Azure**. Tevfik Kosar, State University of New York (SUNY) at Buffalo.
- **Utilizing Continuous Bulk Processing**. Kenneth Yocum, University of California, San Diego.
- **Enabling Mobile Cloud Computing**. Richard Han, University of Colorado Boulder
- **Refining Language Models using Web-scale Language Networks**. Qiaozhu Mei, University of Michigan, Ann Arbor.
- **Predicting Transcription Factor Binding Sites for Genes**. Zhengchang Su, University of North Carolina at Charlotte.
- **Managing Large Watershed Systems**. Jonathan Goodall, University of South Carolina Research Fund and Marty A. Humphrey, the University of Virginia.
- **Tackling Large Scale Graph Problems**. Viktor Prasanna, University of Southern California.
- **Storing Data with Minimal Trust**. Michael Walfish, University of Texas at Austin.
- **Understanding Relational Data Markets**. Magdalena Balazinska, University of Washington.
- **Conducting Intensive Biocomputing**. Wuchun Feng, Virginia Tech.
- **Effectively and Widely Using Renewable Energy Sources**. Kwa-Sur Tam, Virginia Tech.

- **Web Data Analysis Frameworks and Their Applications on Windows Azure**. Tatsuo Nakajima, Waseda University
- **Predicate-Argument Structure Analysis of a Huge Web Corpora for Improving a Search Engine Infrastructure.** Sadao Kurohashi, and a team that includes Daisuke Kawahara, Tomohide Shibata, Keiji Shinzato, and Ryohei Sasano from Kyoto University.
- **Indexing a large Web document collection with modality/factuality information**. Kentaro Inui, Tohoku University.
- **A parallel workflow system for multi-cloud environment.** Kenjiro Taura, University of Tokyo.
- **Ubiquitous Content Management Technology R&D.** Yutaka Kidawara, National Institute of Information and Communications.
- **Inter-Cloud Large-Scale Data Transfer. S**atoshi Matsuoka, Tokyo Institute of Technology.
- **Computing relationships between pairs of vertices in a Wikipedia network, and encrypted and distributed key-value storage on Cloud.** Masatoshi Yoshikawa, Kyoto University.
- **From Content Discovery to Content Distribution - Analysis of Online Social Network Graph**. NICTA. Max Ott, Sebastien Ardon, Anirban Mahanti.
- **Real-time Reasoning for Structured Data on the Web**. NICTA. Kevin Lee project lead. Brief description
- **Computational Chemistry Packages on the Cloud**. NCINF - from the NCI National Facility, Vladislav Vassiliev project lead.
- **High throughput, low impedance eScience on Microsoft Azure.** David Abramson, Monash University.
- **Highly scalable eScience applications**. Raj Buyya, University of Melbourne.
- **Development of a magnetotellurics platform on Azure**. Craig Mudge and Graham Heinson, University of Adelaide.
- From CSIRO a proposed project on management of transportation in the coal industry.

Microsoft