


ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

bioopera

Process Support for Bioinformatics

G. Alonso, W. Bausch, C. Pautasso
Information and Communication Systems Research Group
Department of Computer Science
Swiss Federal Institute of Technology, Zurich (ETHZ)

Who are we?



- Information and Communication Research Group. Our expertise:
 - ◆ High performance, high availability systems (e.g., data replication, backup architectures, parallel data processing)
 - ◆ Information systems architecture (mainframe or cluster based)
 - ◆ Large data repositories (e.g., HEDC data warehouse for the HESSI satellite)
 - ◆ Large distributed systems (e.g., enterprise middleware, workflow, internet electronic commerce)
- Our interests:
 - ◆ System architectures for applications with unconventional requirements
 - ◆ performance: cluster computing and parallel processing
 - ◆ availability: long term fault tolerance, hot backups
 - ◆ communications: ubiquitous computing, mobile ad-hoc networks, internet

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 2

What do we do? **bioopera**

- Until recently, our work has been in the area of high performance/cluster/parallel computing in the IT world
 - ◆ high performance transaction processing systems (throughput, availability)
 - ◆ very large data repositories (scalability, evolution)
 - ◆ data processing applications (large scale data mining)
- In the last years we have started several projects in scientific applications:
 - ◆ HEDC: a multi-terabyte data repository for NASA with the capability to store and process astrophysics data; data production rate is 1 GB per day (raw); data goes on line 24 hours after delivery from the satellite, fully preprocessed, cleaned, catalogued, indexed, with relevant events identified (solar flares, gamma ray bursts, non solar events, etc.), and an extended catalogue of analyses for the events of interest. Users can browse, produce more analyses, request batch processing jobs, perform low resolution analyses, retrieve and store data, etc.
 - ◆ BioOpera: a process based development and run time environment for parallel computations over computer clusters for bio-informatic applications

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

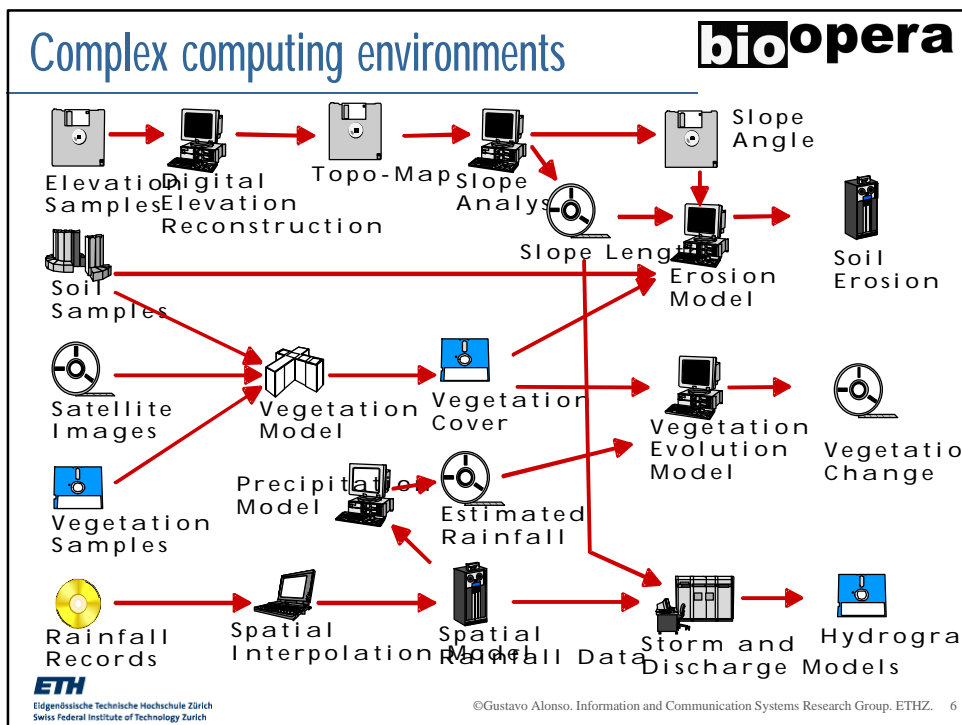
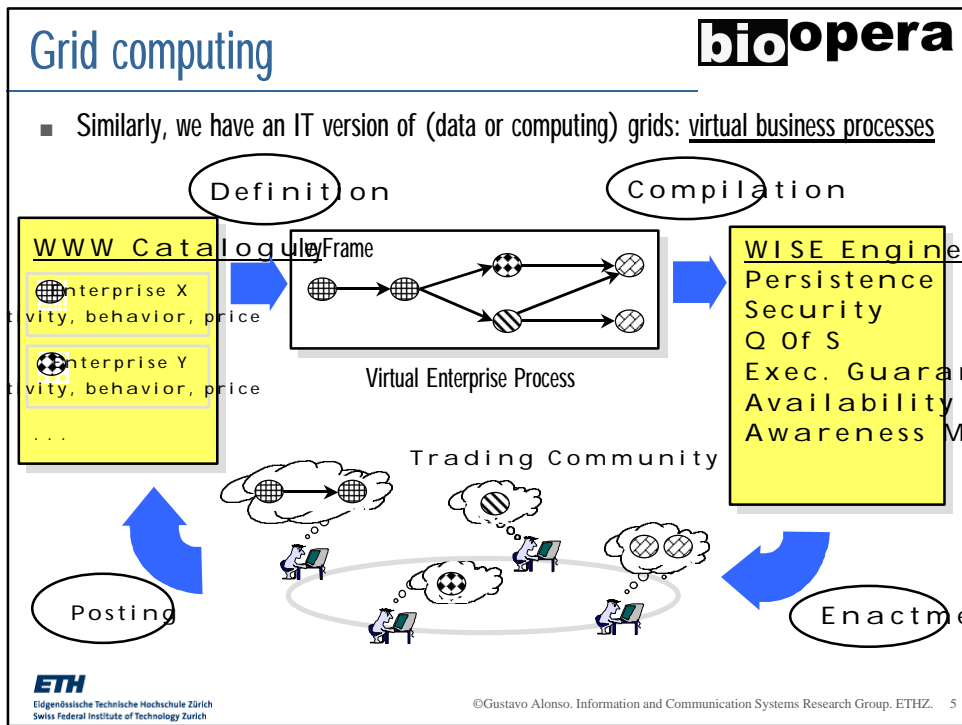
©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 3

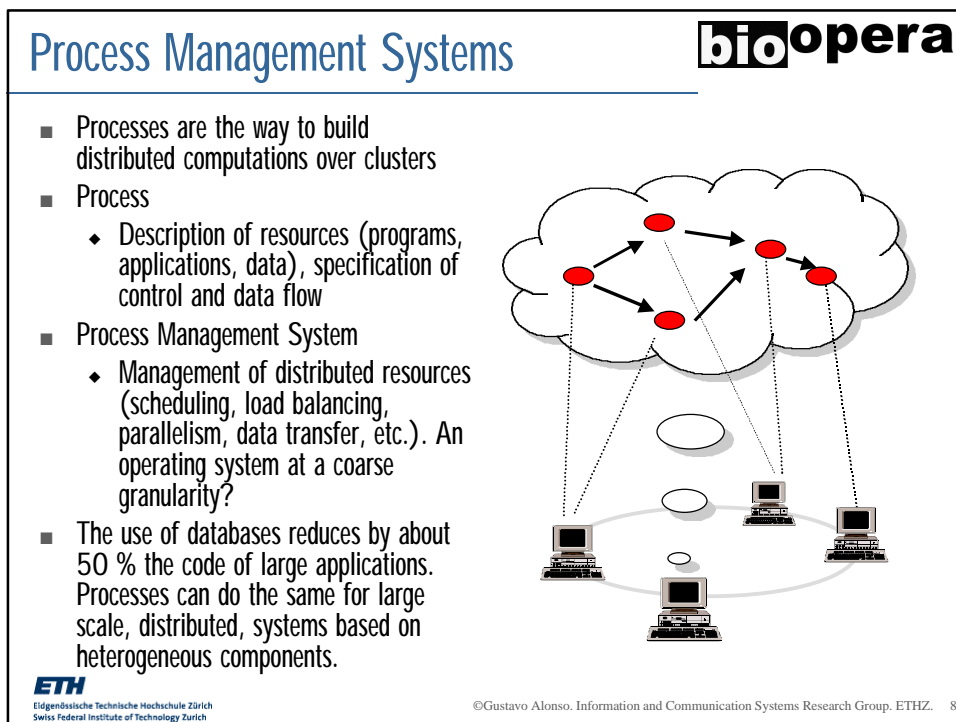
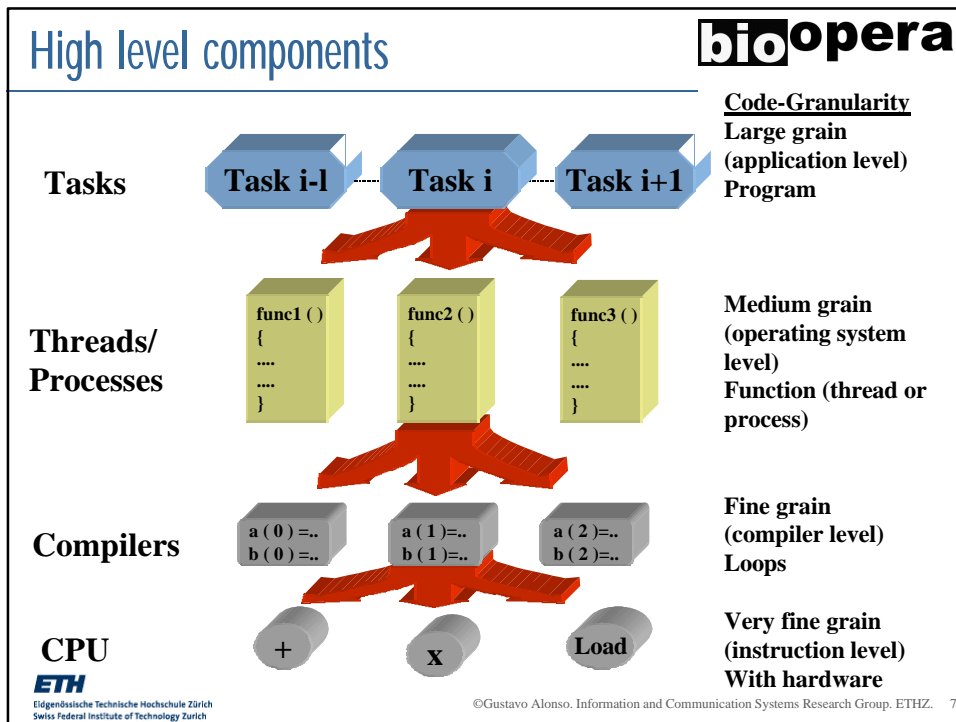
Cluster computing **bioopera**

- Question: "Will cluster computing penetrate the IT market?" (next talk ...)
- Our answer: it did a long time ago but we do not benchmark our clusters in terms of Gflops and there is no Top500 list. Examples:
 - ◆ eBay: cluster of several hundred nodes for high performance transaction processing, fully duplicated for availability purposes
 - ◆ Google: several thousand nodes used for parallel data processing (4 weeks to retrieve, sort, classify, index and store the entire contents of the web)
 - ◆ TPC-C results Compaq ProLiant 8500-700-192P
 - Server: 24 nodes (8 x Pentium III Xeon 700MHz, 8 GB RAM)
 - Clients: 48 nodes (2 x Pentium III 800 Mhz, 521 MB RAM)
 - Disks: > 2500 SCSI Drives
 - Total storage capacity: > 42 TB


ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 4



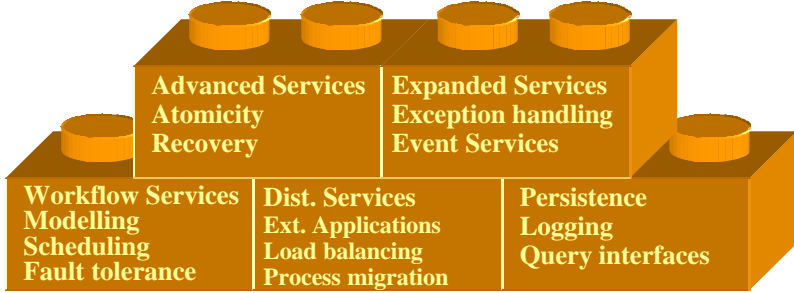


Opera
bioopera



**Open Process Engine
for *Reliable Activities***


- Kernel based system instead of complete application
- Which functionality needs to be implemented in the kernel and which functionality is application specific is an open research area.



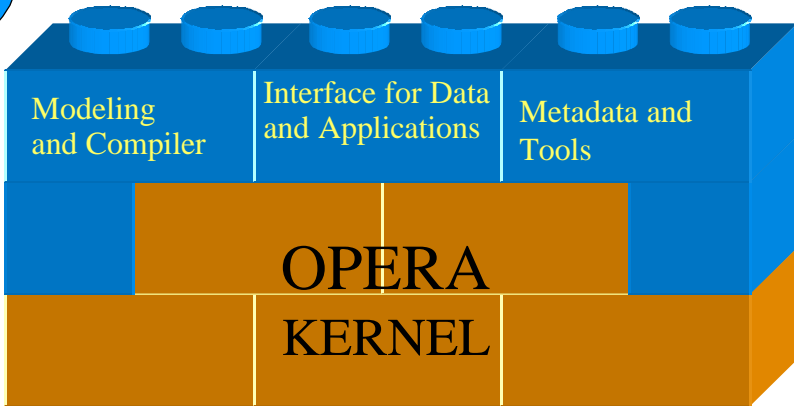
ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 9

Application specific extensions
bioopera



**Open Process Engine
for *Reliable Activities***



ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 10

BioOpera
bioopera

TYPICAL COMPUTING ENVIRONMENTS IN BIO-INFORMATICS

- Multiple step complex computations over heterogeneous applications
 - ◆ algorithmic complexity (most of them NP-hard, NP-complete)
 - ◆ large input / output data sets (exponential is not uncommon)
 - ◆ complex data dependencies and control flow
 - ◆ long lived computations over unreliable platforms
- Current programming environments are very primitive:
 - ◆ support for single algorithms (complex computations put together by hand or very low level scripts)
 - ◆ it is difficult to specify, run, keep track and reproduce computations
 - ◆ no automated support, computations are manually maintained

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 11

BioOpera: A Process Support System
bioopera

BIOPERA

- programming environment for multiple step computations over heterogeneous applications:
 - ◆ program invocations, control & data flow
 - ◆ parallelism (data, search space)
 - ◆ failure behaviour
- Run time environment for processes
 - ◆ navigation
 - ◆ distribution / load balancing
 - ◆ persistence / recovery
- Monitor and interaction tools
 - ◆ at process level
 - ◆ at task level
 - ◆ start, stop, resume, cancel, retry, modify parameters, etc.

Process front-end
 Programming tool
 Monitoring tool

Application front-end
 Data visualization
 Similarity search

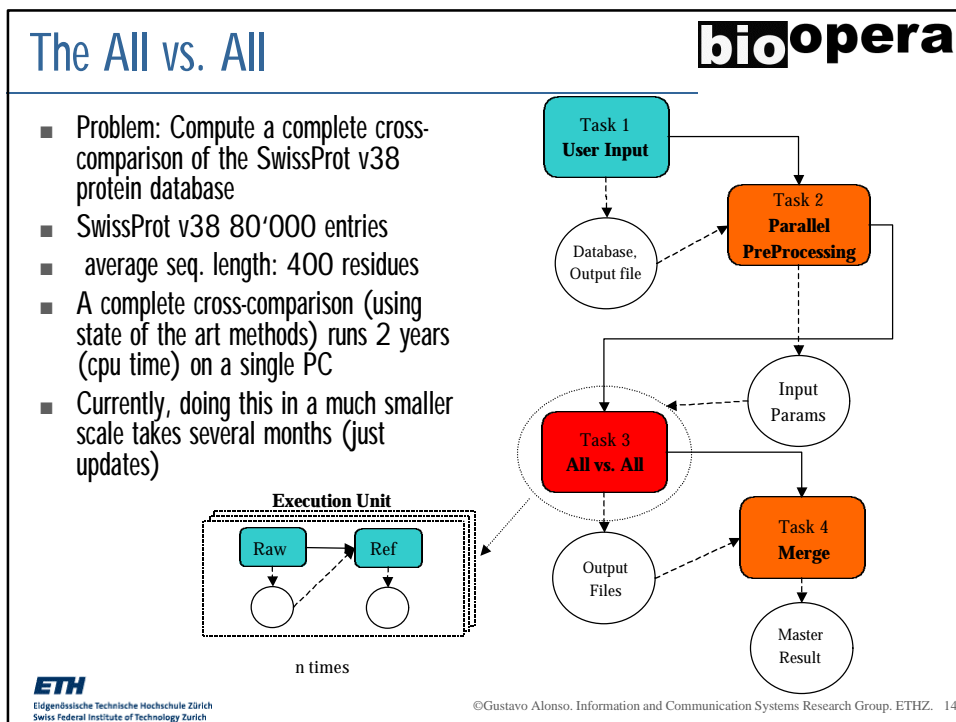
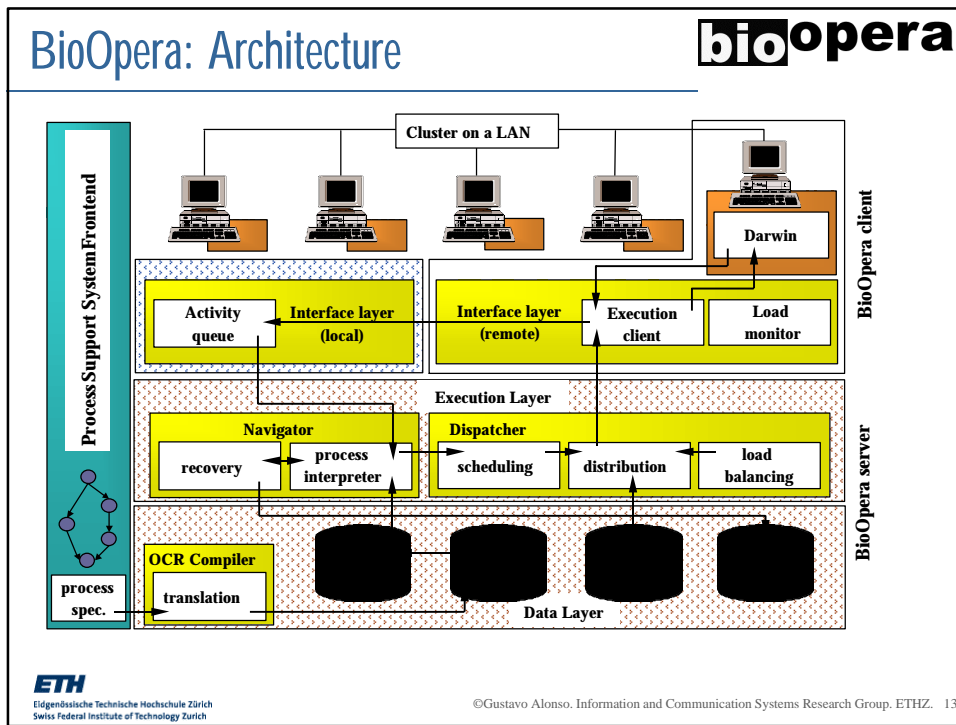
BioOpera
 Process Enactment
 Process Monitoring
 System Awareness

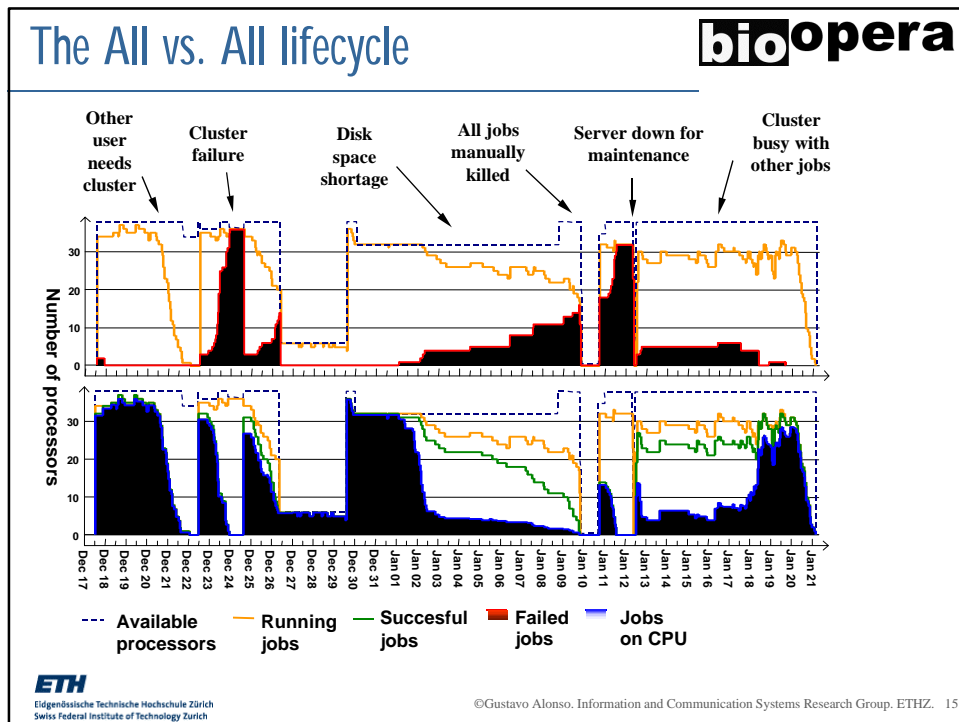
Bioinformatics Application
 DARWIN

Process output

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 12





Conclusions

bioopera

- BioOpera offers:
 - ◆ a better way to program complex computations over clusters,
 - ◆ persistent computations, recoverability,
 - ◆ monitoring and querying facilities,
 - ◆ a scalable architecture that can be applied both in clusters and in grids
- BioOpera is not:
 - ◆ a parallel programming language or a message passing library (too low level),
 - ◆ a scripting language (too inflexible),
 - ◆ a workflow tool (too oriented to human activities)
- BioOpera is extensible:
 - ◆ accepts different process programming models,
 - ◆ monitoring and querying facilities can be expanded and changed,
 - ◆ BioOpera servers can be combined to form clusters of clusters

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

©Gustavo Alonso. Information and Communication Systems Research Group. ETHZ. 16