PillarOne – Scalability of Computing Intensive Models

PillarOne Conference, St. Gallen, 9./10. September 2010

Martin Melchior, PhD, FHNW
Michael Spahn, Intuitive Collaboration AG

Content

- Stochastic simulations applied in insurance
- Distributed stochastic simulations
  - Embarrassingly parallel
  - Requirements and challenges
  - Solution in PillarOne
- Demo and Results

Stochastic simulation techniques (Monte Carlo) used in e.g.

- Modeling low frequency / high severity events (natural perils modeling): Large number of iterations, few simulation periods.
- Multi-period cash flow modeling (e.g. profit testing in life insurance): medium number of iterations, large number of simulation periods (e.g. 40 years, monthly).
- Stochastic simulations in reserving (e.g. bootstrapping): medium number of iterations, medium-large number of simulation periods.

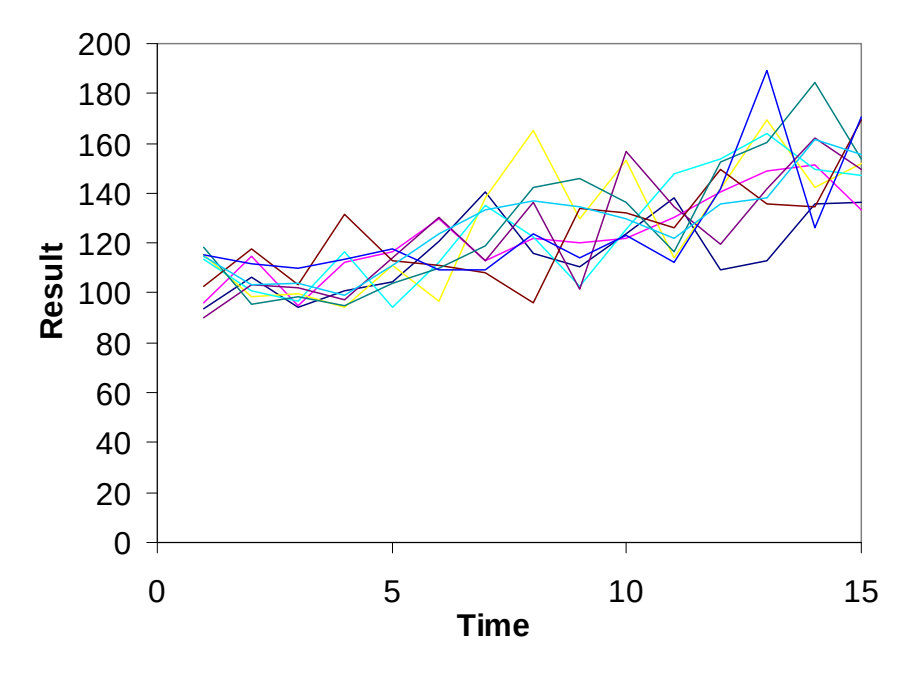

Example: (Partially) Internal **Solvency Models**.

Stochastic Simulations
Example Model in PillarOne

- PODRA model:
  - 4 LoB's
  - Claims:
    - natural perils: storm, earthquake, flood
    - single large, aggregate attritional
  - Simplified modeling of claims development
- Results without distributing the simulations:(*)

Single period, 48'000 iterations

Persistence Solution	Total Time [sec]	Computing	Data Processing
MySQL, Local Files	275	~85%	~15%

(*) **Parameterization: CapitalEagle CorrelatedAttritionals Reinsurance Program NP+ALL50**
Result configuration: CapitalEagle Analysis

Distributed Stochastic Simulations
Embarrassingly Parallel

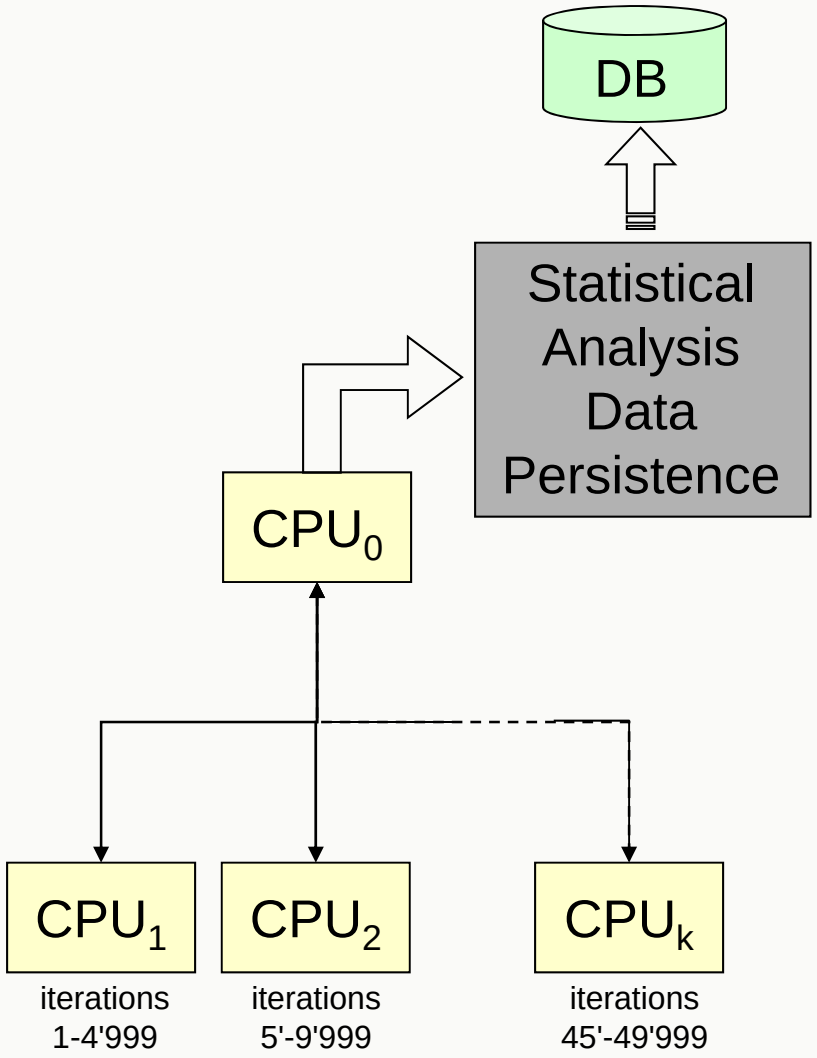

- Generally, in stochastic simulations a large number of **independent** realizations ('iterations') is generated and its output statistically evaluated.
- Independent iterations can be distributed to several independent computing resources (CPU's).
- **Computing** time of the simulations can be largely reduced.
- In general, **data layer not** embarrassingly parallel!

Distributed Stochastic Simulations
Requirements

- Scalability should be a standard feature of the platform: Computing resources are abundant: Even laptops have multi-core processors, servers often have several cpu's, the cloud is waiting.
  PillarOne with its underlying client/server architecture is optimally positioned to be turned into a scalable application.
- Application should allow to easily configure the computing grid. No additional deployment overhead.
- New regulatory requirements demand auditable systems. Results should be reproducible independent of environment and setup – i.e. independent of the computing grid configured.
- Independent of particular model. → Grid configuration a property of the environment and not of the model.

Distributed Stochastic Simulations Challenges

- Computing part is embarrassingly parallel as long as simulated paths on different cpu's are independent. Choosing different seeds for the cpu's does not guarantee independence.
- Guarantee reproducible results independent of grid configuration.
- Grid configuration is a property of the environment and not of the model.
- Data layer not embarrassingly parallel. Iteration data need to be collected for the statistical analysis.
- No limitation on the number of iterations. Data need to be collected before end of simulation – data volume possibly large (level of detail of model, data selected for analysis).

Embarrassingly Parallel Simulations and Reproducible Results

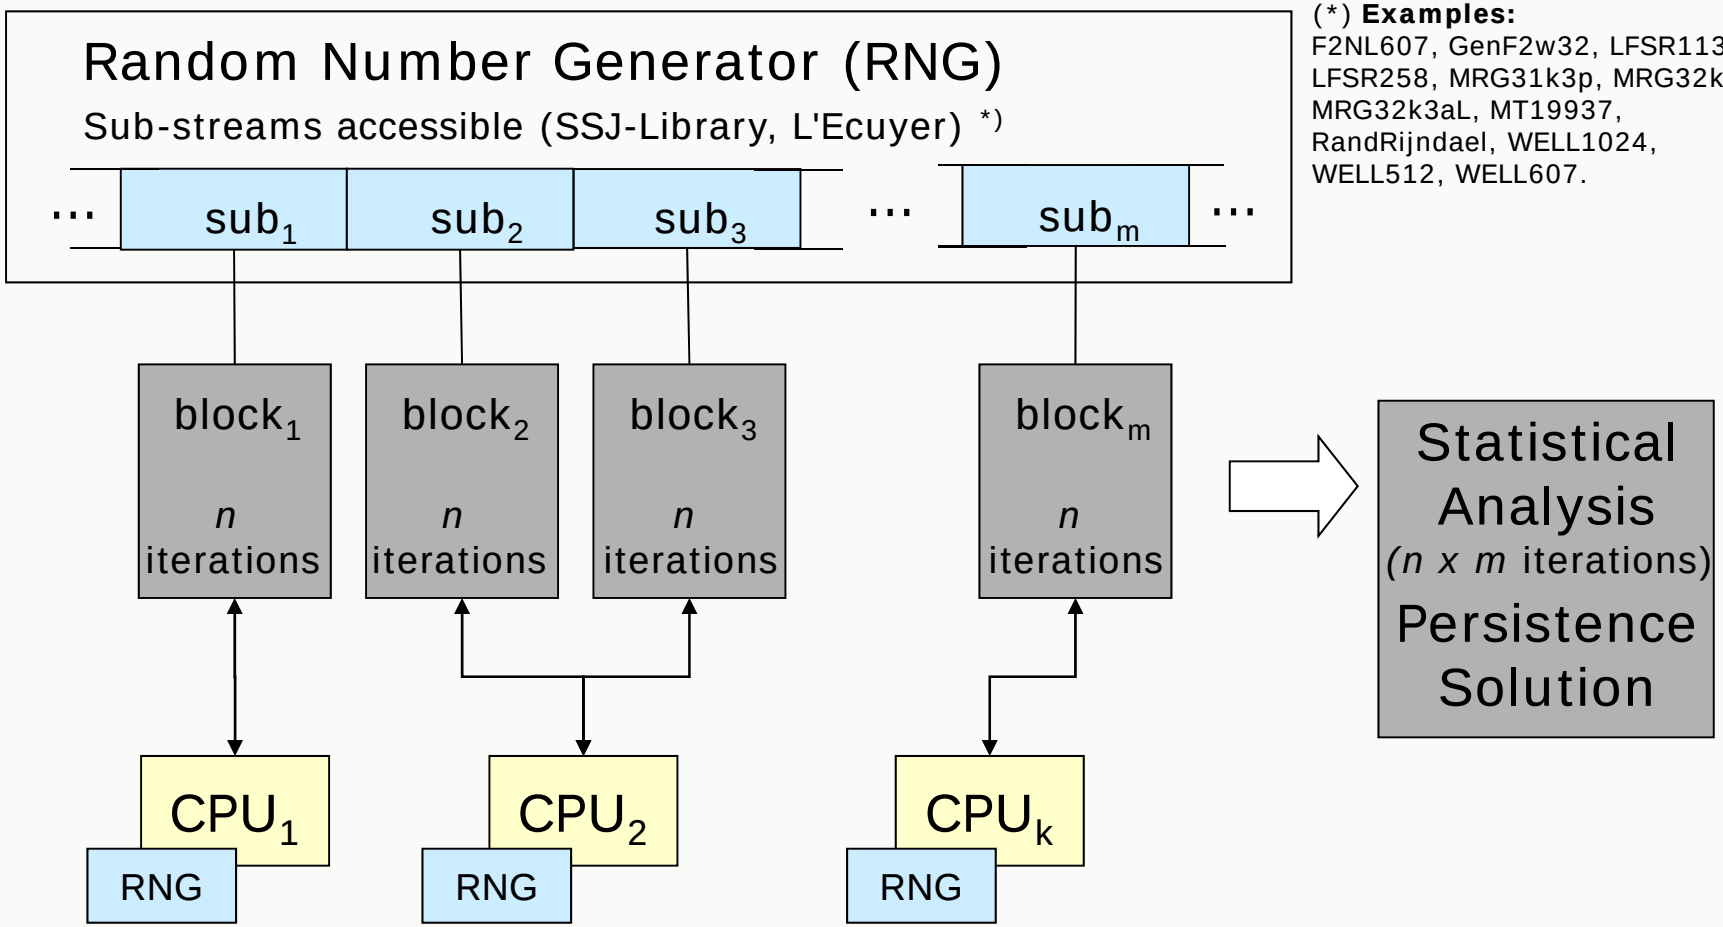

(*) **Examples:**
F2NL607, GenF2w32, LFSR113, LFSR258, MRG31k3p, MRG32k3a, MRG32k3aL, MT19937, RandRijndael, WELL1024, WELL512, WELL607.

java framework *'gridgain'* to distribute simulations.

Data Handling Aspects

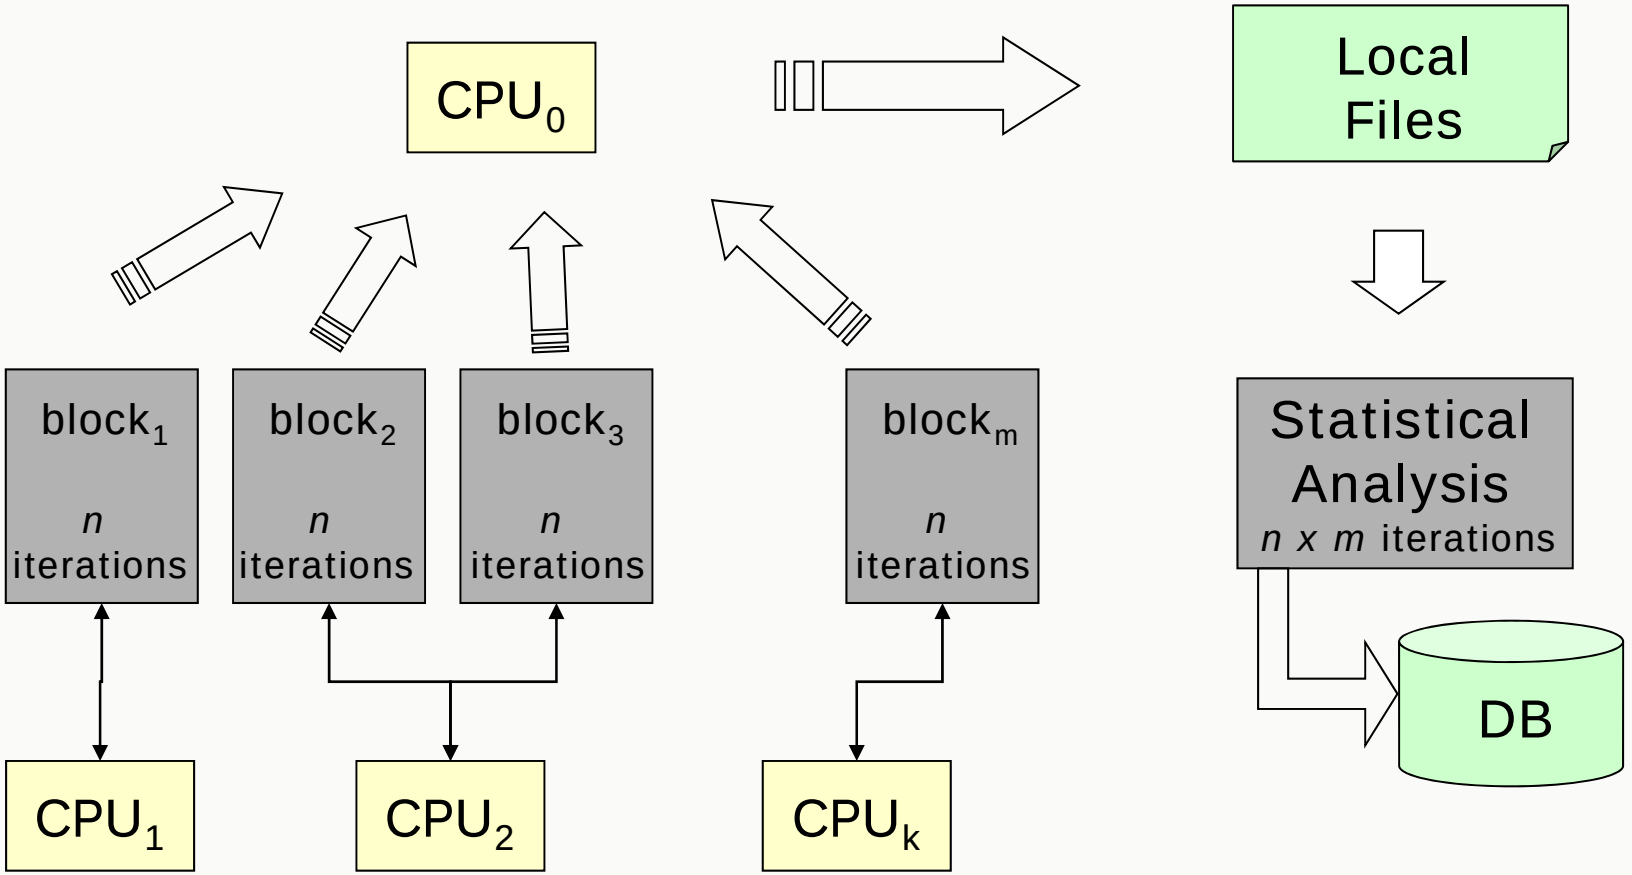

Distributed Stochastic Simulations
Example Model in PillarOne

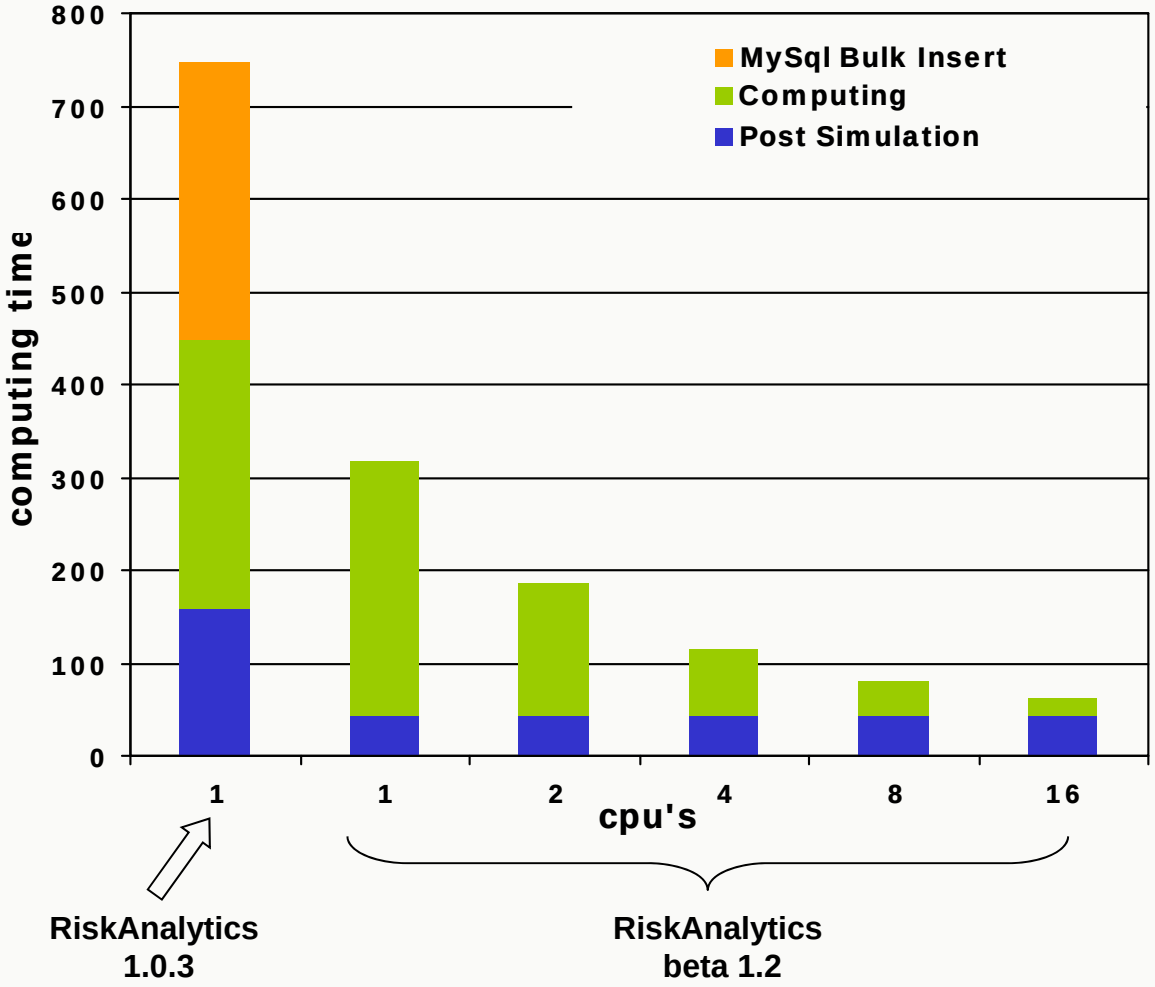

PODRA model
as above,
48' iterations,
15'' records

created with
PillarOne
beta release 1.2

(with plugins:
risk-analytics-core=1.0.3.1.1-kti
risk-analytics-application=1.0.3.1-kti
risk-analytics-pc=1.0.4)

Intel 2.33 GHz

Conclusions

- A framework for distributing simulations has been incorporated into RiskAnalytics. The simulation part proves to be highly scalable.
- The data persistence part is not scalable, but significant improvements in its performance have been implemented.
- In summary, significant improvements in the overall performance for RiskAnalytics are achieved.
- It will be shipped as part of the release RiskAnalytics 1.2.

Contact

Dr. Martin Melchior; Project Manager, Lecturer

+ 41 56 462 40 31; +41 76 561 5616

martin.melchior@fhnw.ch

University for Applied Sciences FHNW, CH-5210 Windisch

Michael Spahn; Software Engineer, Actuarial Tools

+41 44 926 14 08; +41 76 370 31 86

michael.spahn@intuitive-collaboration.com

Intuitive Collaboration AG; Seestrasse 16; CH-8712 Staefa

Backup

Moore's Law

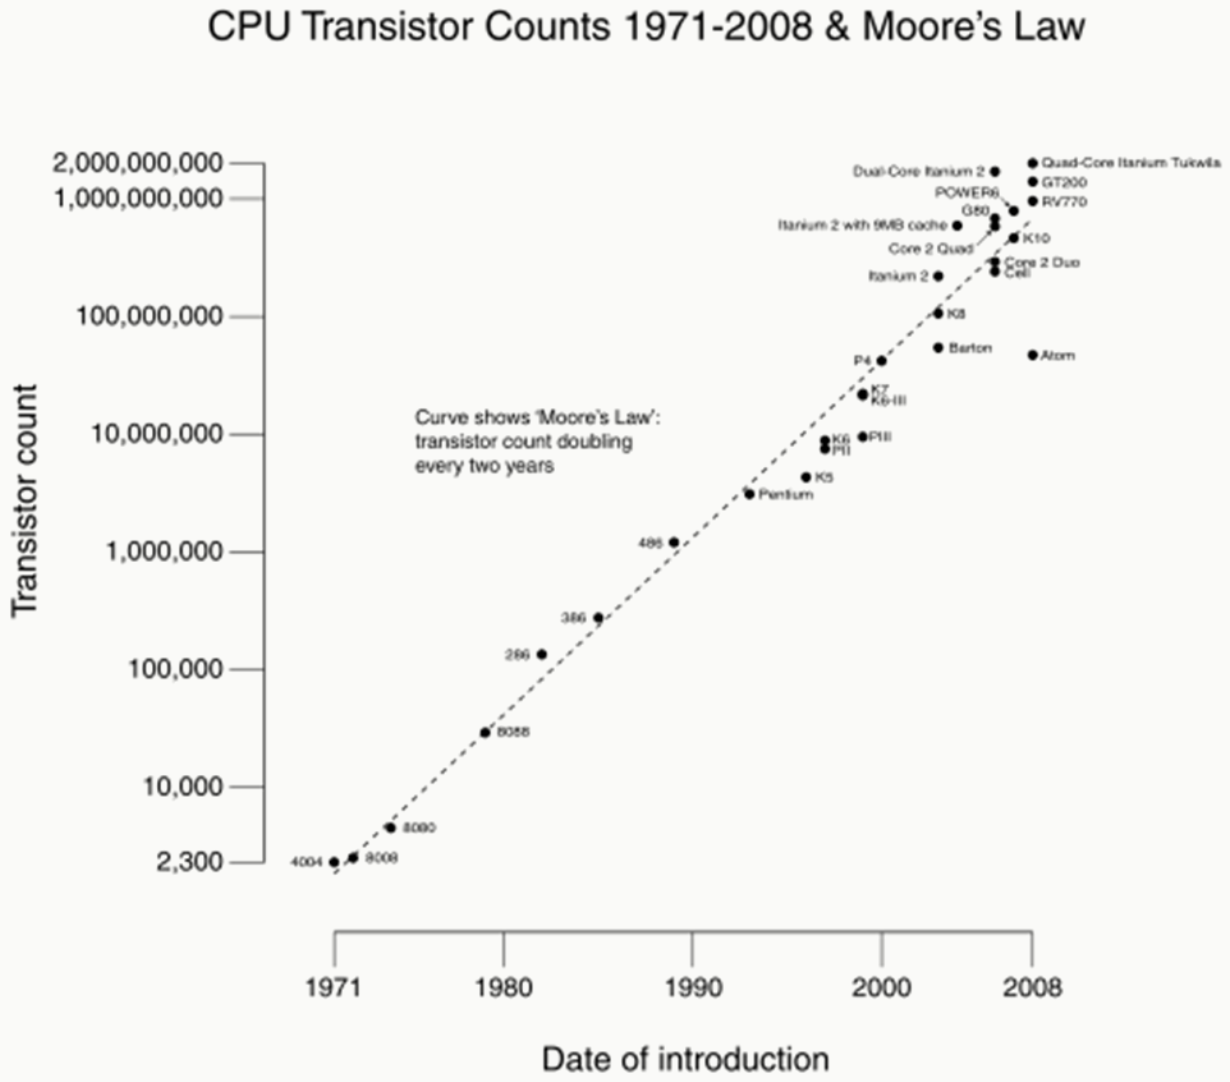
CPU Transistor Counts 1971-2008 & Moore's Law
Transistor count
2,000,000,000
1,000,000,000
100,000,000
10,000,000
1,000,000
100,000
10,000
2,300
Curve shows 'Moore's Law': transistor count doubling every two years
Dual-Core Itanium 2
Quad-Core Itanium Tukwila
GT200
POWER6
RV770
G80
Itanium 2 with 9MB cache
K10
Core 2 Quad
Core 2 Duo
Cell
Itanium 2
K8
Barton
Atom
P4
K7
K6-III
K6
P II
P III
K5
Pentium
486
386
286
8088
8080
4004
8008
1971
1980
1990
2000
2008
Date of introduction

CTI Project: Turning PillarOne Plattform in a Scalable Application

- Joint project between Intuitive Collaboration AG and FHNW: "*Risk Management Platform PillarOne - safe model building, efficient computing, guided analysis".* FHNW contribution is to a large extent funded by the Swiss Government, Commission for Technology and Innovation (CTI).
- Efficient computing:
  - Allow distributing simulations to many CPU's, processors, computers or computing servers ...
  - ... such that RiskAnalytics becomes *scalable*: Providing a factor 10 more computing resources should reduce the runtime by nearly a factor 10.

Distributed Stochastic Simulations
Example Model in PillarOne

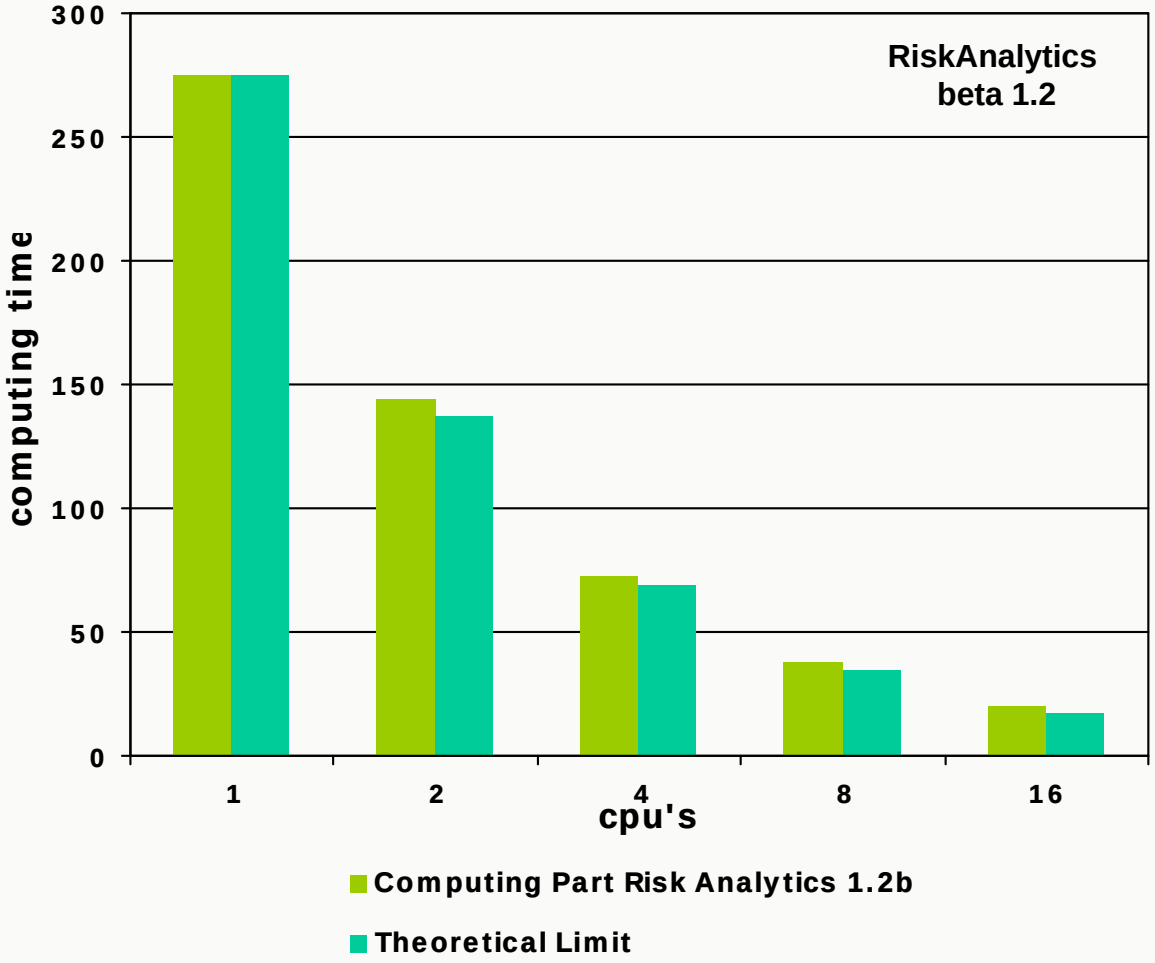

PODRA model
as above,
48' iterations,
15'' records

created with
PillarOne
beta release 1.2

(with plugins:
risk-analytics-core=1.0.3.1.1-kti
risk-analytics-application=1.0.3.1-kti
risk-analytics-pc=1.0.4)

Intel 2.33 GHz