Tutorial at AICCSA'07
(May 13-16, 2007, Amman, Jordan)

Title: **Introduction to Accurate Stochastic Discrete-Event Simulation**

Duration:
a half day tutorial

Instructor:
Krzysztof (Krys) Pawlikowski
Professor in the Department of Computer Science and Software Engineering
University of Canterbury in Christchurch, New Zealand

Summary:

The computer revolution initiated in the twentieth century has resulted in the adoption of computer simulation as the most popular paradigm of scientific and engineering investigations. It has become the most commonly used tool in performance evaluation studies of such complex dynamic stochastic systems as telecommunication networks. Such reliance on simulation studies raises the question of credibility of the results they yield.

In this tutorial, we will discuss **main problems and solutions of quantitative stochastic discrete-event simulation**, i.e. the stochastic simulation in which the emphasis is put on statistical correctness of the final results. Whole spectrum of the problems will be covered: from generators of pseudo-random numbers to methods of controlling precision of the final results in sequential stochastic simulation conducted in a distributed mode on multiple computers of a LAN. The latter will be discussed in the context of Multiple Replications in Parallel (MRIP) scenario, which allows to speed up stochastic simulation by launching multiple simulation engines cooperating in production of data for central analysers (one central analyser for each performance measure). An implementation of MRIP in a simulation package Akaroa2, a universal controller of distributed quantitative simulation, which automatically launches simulations on an arbitrary number of simulation engines and controls precision of an arbitrary number of analysed performance measures, both in terminating and steady-state simulation, will be also discussed.

The tutorial will conclude with a survey of open research problems of quantitative stochastic simulation

Outline of Topics:
The topics will be discussed in the following order:

- Introduction to quantitative stochastic simulation: basic terms and definitions. Necessary conditions of a trustworthy simulation study.
- Sources of randomness: theory and practise of modern generators of pseudo-random numbers for simulation on single and multiple processors.
- Sequential quantitative stochastic simulation: its principles and implementation in terminating, steady-state, and non-stationary simulation; measures of precision of the final results.
- Automation of precision control in terminating and steady-state simulation: survey of methods of analysis. The initial transient period in steady-state simulation: theory and methods for estimating its length.

- Multiple Replications in Parallel (MRIP) for speeding up quantitative stochastic simulation. Applications of MRIP in Akaroa.2.

Additional readings:

- o M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, 2006.
- o K. Pawlikowski. Steady-state Simulation of Queueing Processes: A Survey of Problems and Solutions. *ACM Computing Surveys,* June 1990, 123-170

Prerequisities.
It is assumed that participants will be familiar only with the very basic concepts of probability theory and statistics.

Biography of the instructor:
Speaker Biography
Prof. Krzysztof (Krys) Pawlikowski has over 30 years experience as a university lecturer. Currently he is a Professor in Computer Science and Engineering at the University of Canterbury, in Christchurch, New Zealand.

He received his PhD degree in Computer Engineering from the Technical University of Gdansk, Poland. The author of over 140 research papers and four books, has given invited lectures at international conferences and at over 80 universities and research institutes in Asia, Australia, Europe and North America.

His research interests include performance modelling of multimedia telecommunication networks, tele-traffic modelling, computer simulation and distributed processing.

Target Audience:
Research students, scientists and engineers who use discrete-event simulation in performance evaluation studies of computer systems and networks, or any other stochastic dynamic systems, such as communication or inventory systems, production lines, etc.

Introduction to Accurate Stochastic Discrete-Event Simulation

Krzysztof Pawlikowski

University of Canterbury
Christchurch


Kaitaia
Kerikeri
Paihia
Whangarei
Dargaville
Warkworth
Helensville
AUCKLAND
Coromandel
Thames
Paeroa
Waihi
Katikati
Tauranga
Te Puke
Hicks Bay
Whakatane
Opotiki
NORTH ISLAND
Hamilton
Otorohanga
Rotorua
Tokoroa
Taupo
Waikaremoana
Gisborne
New Plymouth
Turangi
Wairoa
Waiouru
Napier
Hawera
Wanganui
Palmerston North
Woodville
Paraparaumu
Masterton
Collingwood
WELLINGTON
Motueka
Picton
Nelson
Blenheim
Westport
Reefton
Punakaiki
Springs Junction
Kaikoura
Greymouth
Lewis Pass
Hanmer Springs
Hokitika
Arthurs Pass
Waipara
Franz Josef
CHRISTCHURCH
Mount Cook
Haast
Ashburton
Akaroa
Timaru
Omarama
SOUTH ISLAND
Milford Sound
Wanaka
Kurow
Queenstown
Cromwell
Oamaru
Alexandra
Te Anau
Manapouri
DUNEDIN
Gore
Milton
Invercargill
Half Moon Bay
STEWART ISLAND

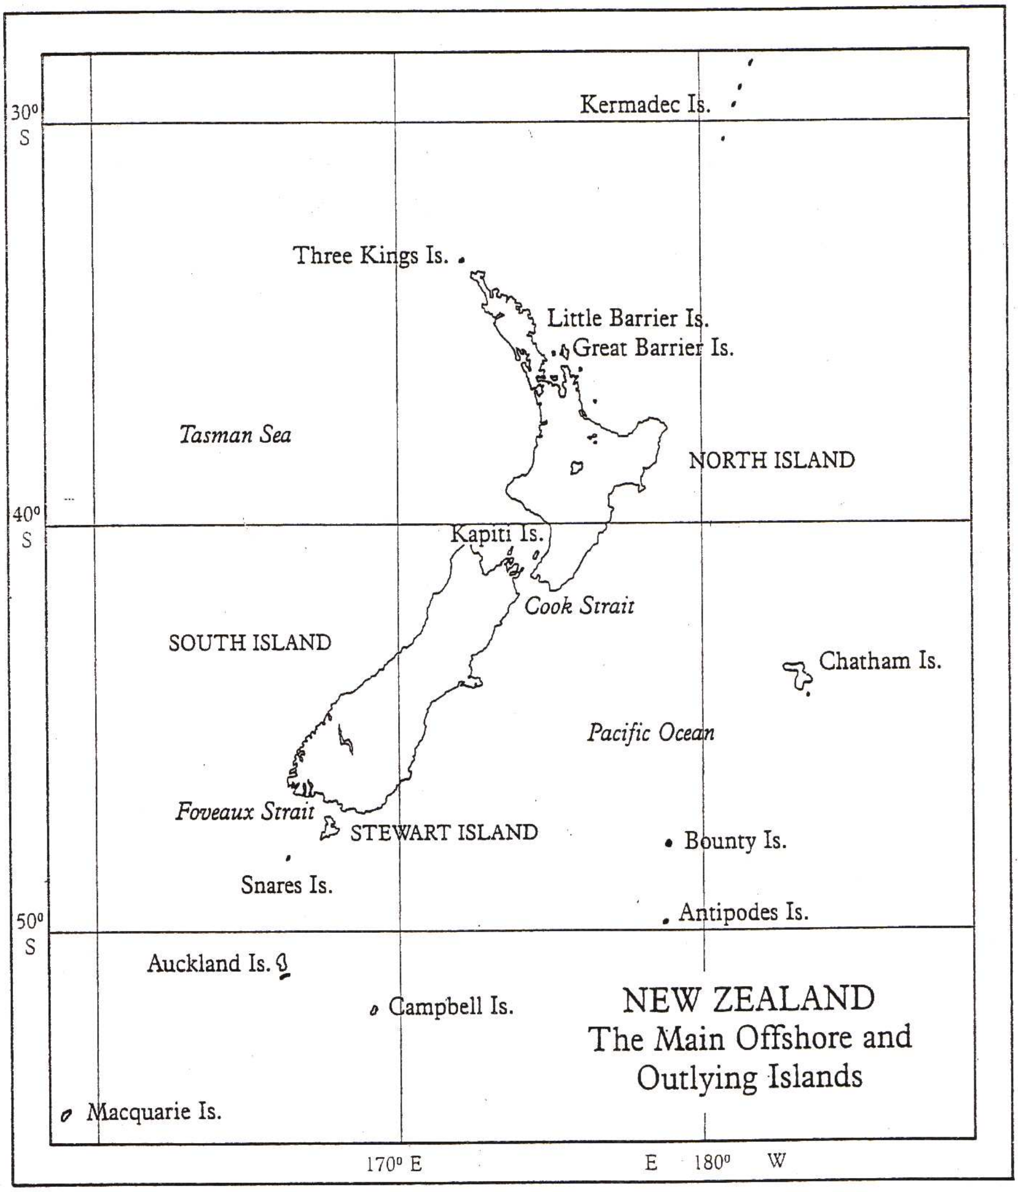
Kermadec Is.
Three Kings Is.
Little Barrier Is.
Great Barrier Is.
Tasman Sea
NORTH ISLAND
Kapiti Is.
Cook Strait
SOUTH ISLAND
Chatham Is.
Pacific Ocean
Foveaux Strait
STEWART ISLAND
Bounty Is.
Snares Is.
Antipodes Is.
Auckland Is.
Campbell Is.
NEW ZEALAND
The Main Offshore and
Outlying Islands
Macquarie Is.
30° S
40° S
50° S
170° E
E 180° W

University of Canterbury: Interdepartmental Network Research Laboratory:

Protocols, Distributed Processing and Simulation

(1) Simulation Research Group:
developing new methodology of distributed and automated stochastic discrete-event simulation

(2) Network Research Group:
modelling and evaluating the Internet
and other multimedia networks
in wired, wireless and optical technologies

Membership:

- Prof. Dr K. Pawlikowski (CSSE)
- Prof. Dr H. Sirisena (ECE)
- Associate Prof. Dr D. McNickle (MGMT)
- Dr. Greg Ewing (research associate)
- Master and PhD students of CSSE, ECE and Management

In 2000-2005:
5 PhD graduates (3 from CSSE)
over 10 Master graduates (4 from CSSE)

Currently: 4 PhD and 6 Master students
(1 PhD and 2 MSc students of CSSE)

see www.cosc.canterbury.ac.nz/research/RG/net_sim/

Computer simulation:

one of the most important factors enabling new developments in science & technology

Congress of the USA, 1996

Three paradigms of science:

- theory
- experimentation
- computer simulation

Interdisciplinary character of Computer Simulation

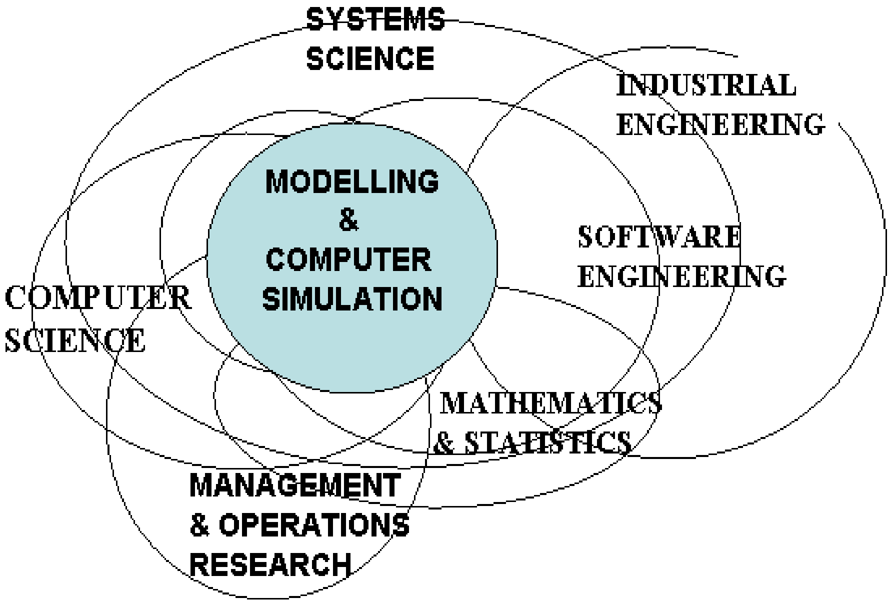

Current controversy:
Should **Computer Simulation** be considered as a new discipline (of Computing) ?
See discussion within the Society for Modelling and Computer Simulation.

Departments of (Modeling and) Computer Simulation *have been established at a number of universities world-wide ...*

Quantitative stochastic simulation:

used in performance evaluation studies of dynamic stochastic systems and processes

In particular, studies of
Telecommunication Networks, motivated by:

- Enormous dependence of our civilisation on computers and their networks makes their performance evaluation mandatory before they are deployed.
- Modern multi-media networks have become so complex that their performance can be usually studied only by simulation.

Advanced Research Project Agency of US Department of Defense lists **Network Modelling & Simulation** *among the most important 23 areas of Information Technology which they support.*

"DARPA Information Processing Technology Office will create a new generation of computational and information systems that possess capabilities far beyond those of current systems".
For this purpose 23 IT programmes have been identified, including:

Architectures for Cognitive Information Processing
Bio-Computation
Compact Aids for Speech Translation
Coordination Decision Support Assistants
Effective, Affordable, Reusable Speech-to-Text
Fast Connectivity for Coalitions and Agents Project
High Productivity Computing Systems
Learning Applied to Ground Robots
Mobile Autonomous Robot Software
Network Modeling and Simulation
Organically Assured and Survivable Information Systems
Personalized Assistant that Learns
Polymorphous Computing Architectures
Power Aware Computing and Communication
Quantum Information Science and Technology
Real-World Reasoning
Self-Regenerative Systems
Software for Distributed Robotics
Taskable Agent Software Kit

see *www.darpa.mil/ipto/index.htm* (in July 2005)

Challenge:
to find solutions
(protocols, architectures, topologies)
which can ensure the best (or the worst) performance of networks

Issue of credibility of simulation studies:

- When can results obtained by quantitative stochastic simulation be regarded as **accurate** ?
- How to ensure **credibility** of quantitative analysis based on stochastic simulation ?

General guidelines for conducting valid performance evaluation studies based on stochastic simulation:

- use a valid simulation model:
  - correct functional representation of simulated mechanisms
  - appropriate stochastic characteristics of the model
- **execute valid simulation experiment:**
  - **use appropriate pseudo-random number generators (PRNGs)**
  - **apply appropriate analysis of simulation output data**

SOURCE(S) OF PRIMARY RANDOMNESS:

Good PRNGs (of uniformly and independently distributed numbers):

- should generate numbers that satisfy the most rigorous statistical tests (of uniformity and independency),
- should be able to generate multiple streams of numbers in long cycles,
- should be fast and accurate.

Until now, the most popular PRNGs are linear congruential generators
with cycle $L \leq 2^{48} - 1$
(with $L \leq 2^{31} - 1$ being still the most popular)

Taking into account recent developments in science and technology, one can show these "baby" generators cannot be considered as good PRNGs and should be replaced by much more powerful ones.

Cycle length vs. computing technology

Proposition:

The length of a typical simulation (measured by CPU time) does not become shorter as computing technology becomes more powerful: **users use faster computers for simulating more complicated processes within the available time**.

Conclusion 1:

PRNGs with adequately long cycles are needed to avoid repetitions of Pseudo-Random Numbers (PRNs) during a typical simulation.

Technology & generation rate of PRNs

Moore's law by Gordon Moore (1965)
see *research.microsoft.com/rgray/Moore_Law.html*

Law of accelerating returns by Ray Kurzweil (2001)
see *www.kurzweilai.net/articles/art0134.html*

On average, computing power of computers doubles each 1.5 - 2 years

At University of Canterbury in Christchurch, New Zealand:

In 1995, workstations with CPU clock at 150 MHz
using rand (from C library): 10^6 PRNs in about 2 seconds, or
2^{31} PRNs in about 72 minutes
2^{48} PRNs in about 18 years

In 2000, PCs with the CPU clock at 800 MHz
using rand (from C library): 10^6 PRNs in about 0.4 seconds, or
2^{31} PRNs in about 13.5 minutes
2^{48} PRNs in about 3.5 years

In 2002, PCs with the CPU clock at 2.4 GHz
using rand (from C library): 10^6 PRNs in 0.13 seconds, or
2^{31} PRNs in about 4.5 minutes

2^{48} PRNs in less than 1.1 years

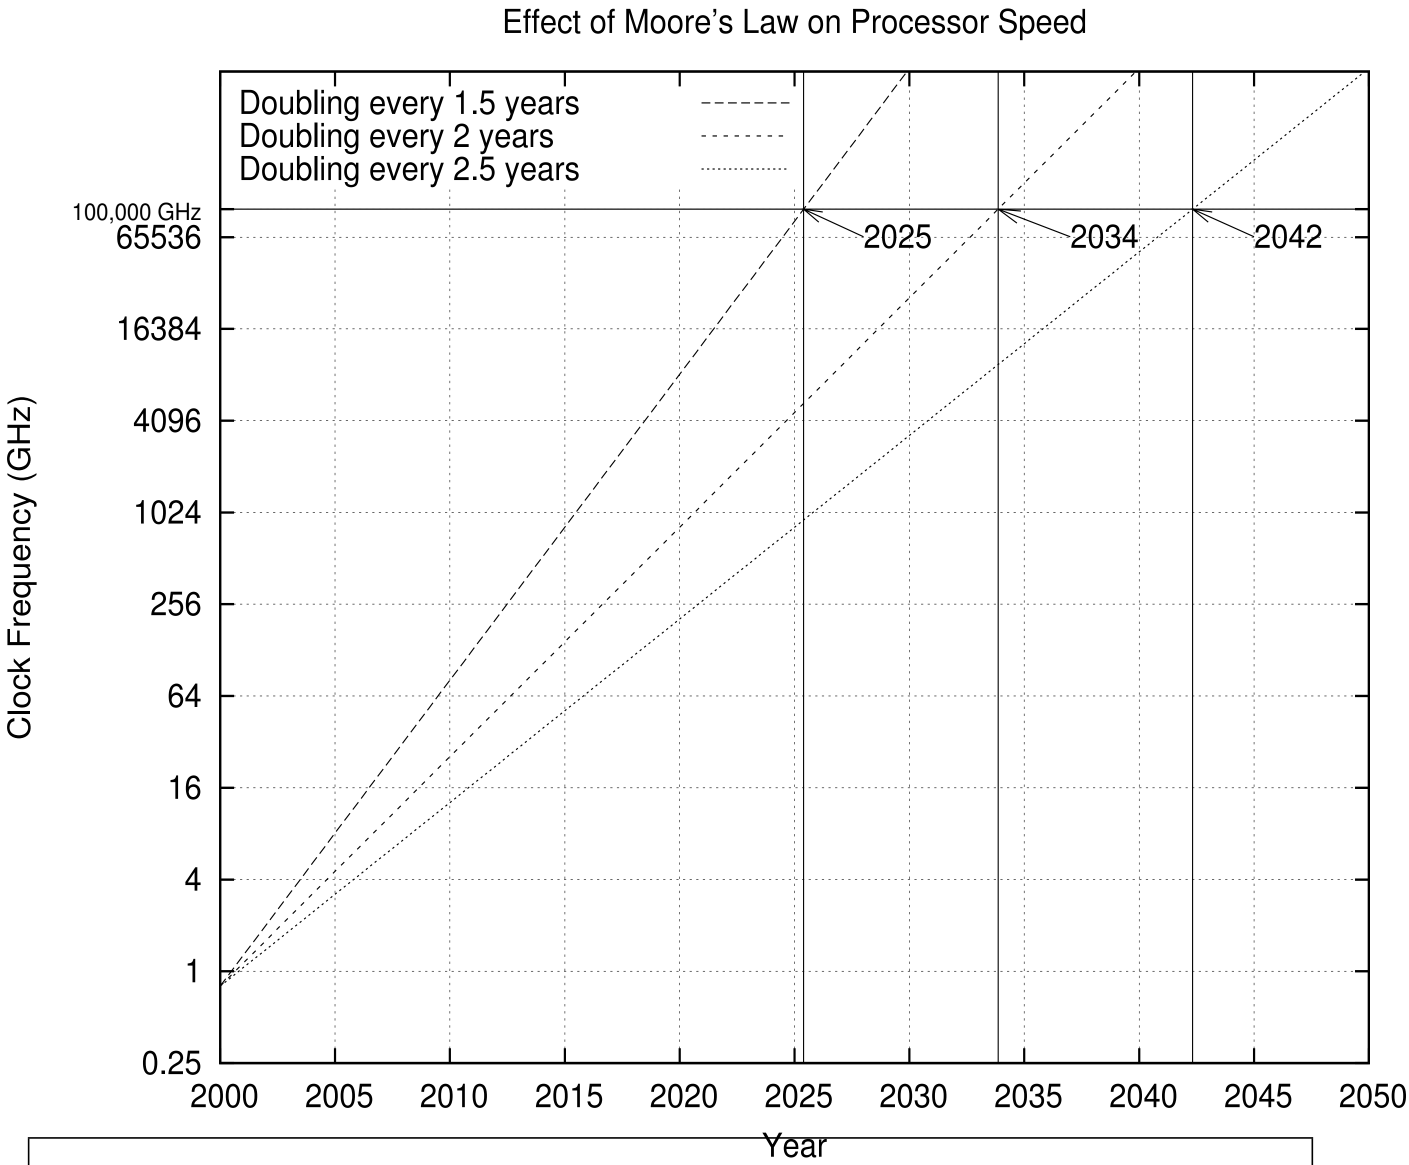

If clock frequency continues to double each 1.5 (or 2, or 2.5 years), then we will have CPU running at over 100 THz in 2025 (or in 2034, or in 2042).

Cycle length of PRNGs vs. theory

Birthday Spacing Problem:
one of reference problems/tests for testing uniformity of PRNGs
see D.E.Knuth. Art of Computer Programming, vol.2

Recent finding:

Any linear congruential PRNG fails the Birthday Spacing Test if one applies this test to $n \geq 16\sqrt[3]{L}$ numbers generated by that PRNG, where L is the length of its cycle.

L'Ecuyer and Simard;
in Mathematics and Computers in Simulation, 2001.

Conclusion 2:

During a simulation,
a PRNG should be used as a source of **no more than** $16\sqrt[3]{L}$ **numbers** (linear congruential PRNG), or **no more than L numbers** (non-linear congruential PRNG).

Using popular linear congruential PRNGs on computers with 2.4 GHz CPU clock

Cycle length	n_{max}	**Max. total time of “stochastically safe” application of a good PRNG**
$L = 2^{31}$	20 594	less than 0.3 seconds
$L = 2^{48}$	1 048 576	less than 14 seconds

where $n_{max} = 16\sqrt[3]{L}$

Search for good PRNGs of the 21^{st} century

M. Schoo, K. Pawlikowski & D. McNickle,
Technical Report TR-COSC 03/05

Assumption:

the process of PRN generation takes no longer than 1% of simulation time.

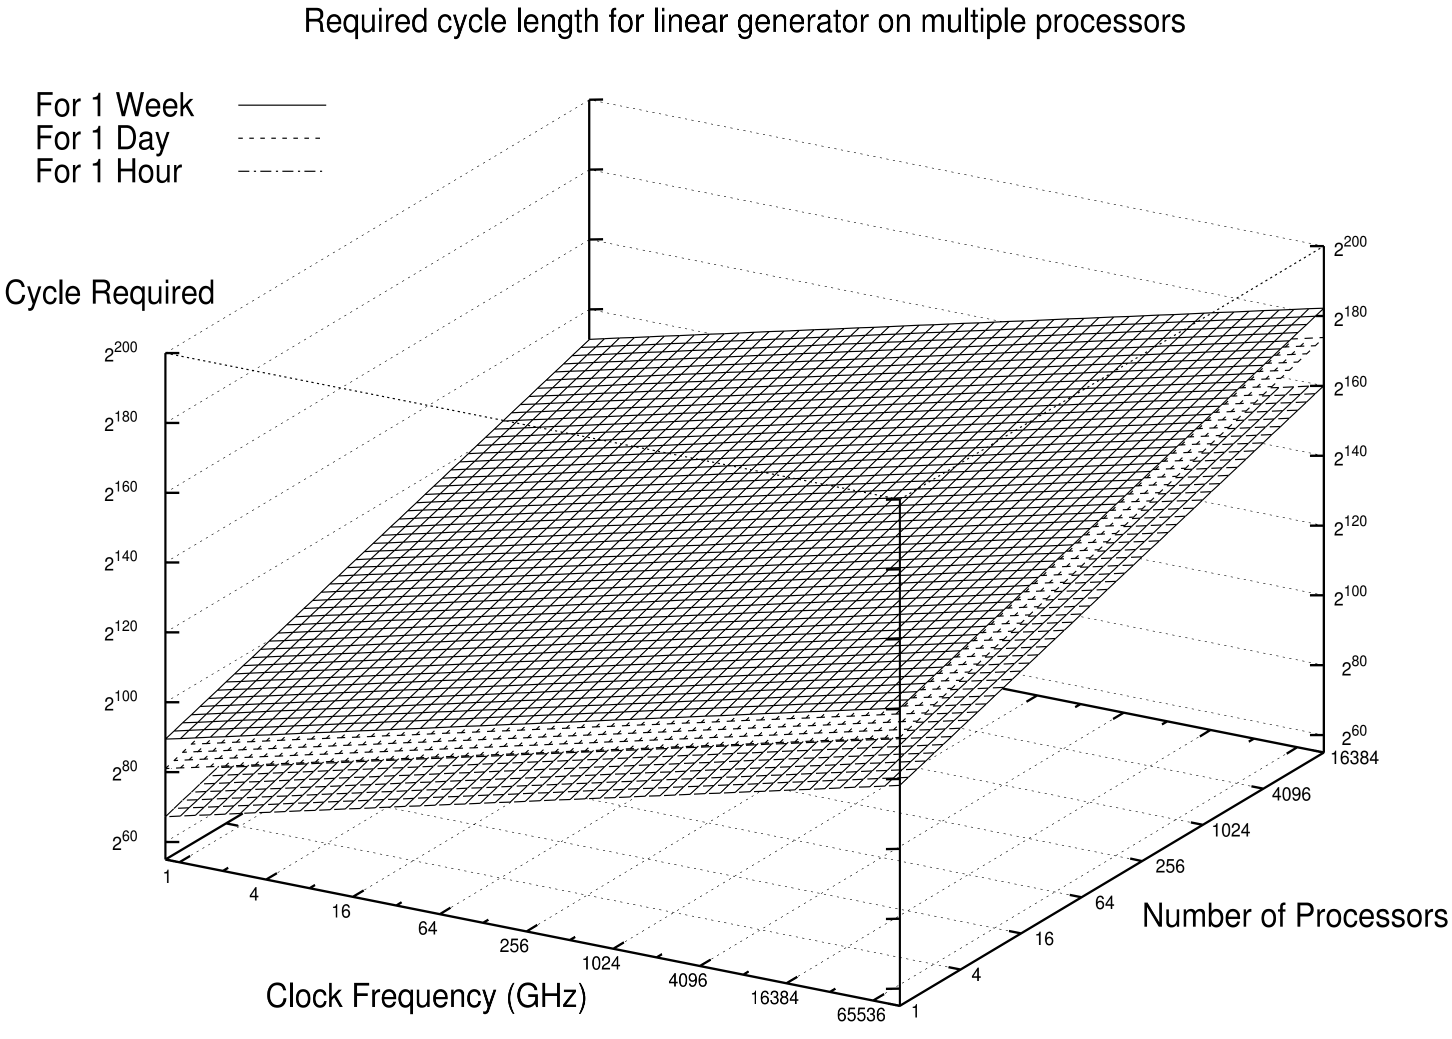

Distributed simulation on 2^{14} processors, with 100 THz CPUs, needs a linear congruential PRNG with the cycle length of about 2^{160} if it lasts 1 hour, and the cycle length of about 2^{182} if it lasts 1 day.

Required cycle length for non-linear generator on multiple processors

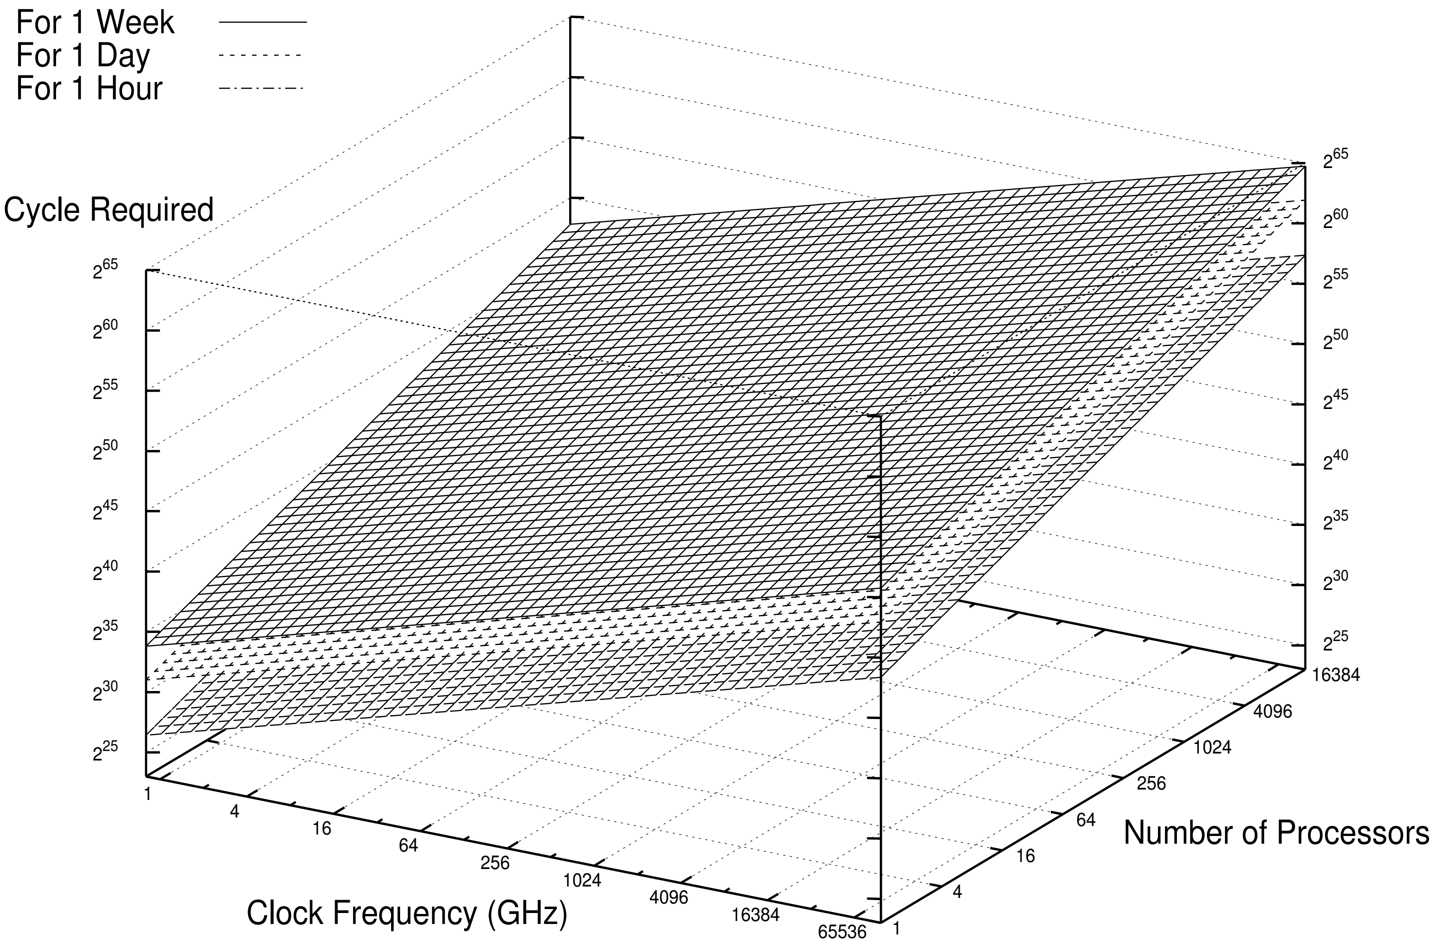

Distributed simulation on 2^{14} processors, with 100 THz CPUs, needs a nonlinear congruential PRNG with the cycle length of about 2^{57} if it lasts 1 hour, and the cycle length of about 2^{65} if it lasts 1 day.

Conclusion:

"Statistically safe"' PRNGs
of the 21^{st} century should be able to generate numbers in cycles of length $L > 2^{180}$ (linear PRNGs), or $L > 2^{65}$ (non-linear PRNGs).

PRNGs of the 21st century:

Multiple Recursive Congruential PRNGs:

G1:

$$r_i = (2109532706 r_{i-1} + 1651737654 r_{i-7}) \bmod (2^{31} - 19)$$

$$L = (2^{31} - 19)^7 - 1 \approx 2^{217}$$

Multiple streams of PRNs generated by **Cycle Splitting**. Proven good uniformity in up to 20 dimensions.

For other properties and portable implementations, see: *L'Ecuyer, Blouin and Couture. ACM Trans. on Modelling and Computer Simulation, April 1993*.

Combined Multiple Recursive Congruential PRNGs:

G2:

$$r_{1,i} = (1403580 r_{1,i-2} - 810728 r_{1,i-3}) \bmod (2^{32} - 209)$$
$$r_{2,i} = (527612 r_{2,i-1} - 1370589 r_{2,i-3}) \bmod (2^{32} - 22853)$$
$$ri = (r_{1,i} - r_{2,i}) mod (2^{32} - 209)$$
$$L \approx 2^{191}$$

Multiple streams of PRNs generated by **Cycle Splitting**. Proven good uniformity in up to 45 dimensions.

For other properties and portable implementation of G2, see: *P. L'Ecuyer. Operations Research, 1996*.

Generalised Feedback Shift Register (GFSR) PRNGs:

Twisted GFSR PRNG known as Mersenne Twister (MT19937)
with the cycle $L = 2^{19937} - 1$

M. Matsumoto and T. Nishimura. ACM Trans. Modelling and Computer Simulation, 1998

Some properties of MT19937:

- uniformity in up to 623 dimensions;
- multiple streams of PRNs generated by **dynamic creation of multiple Mersenne Twisters**

Warning: *There is a problem with initialization. Special care has to be taken when selecting the seed, as a long sequence of strongly correlated numbers can be generated !*

Note:

Non-linear PRNGs (Inversive Congruential and Explicite Inversive Congruential) are currently too slow to be practical in real applications. A fast algorithm for modular inversion is needed.

Our recommendation:

- Use one of linear congruential PRNGs proposed by Pierre L'Ecuyer.

- Multiplicative Lagged-Fibonacci PRNGs can be also considered.

SIMULATION OUTPUT DATA ANALYSIS:

stochastic simulation	$\equiv$	simulated statistical experiment

Results of any quantitative simulation study can be credible only if they are obtained with an acceptable (low) statistical errors.

Otherwise ... the results can be meaningless !

Which of these two sets of results are more conclusive/credible ?

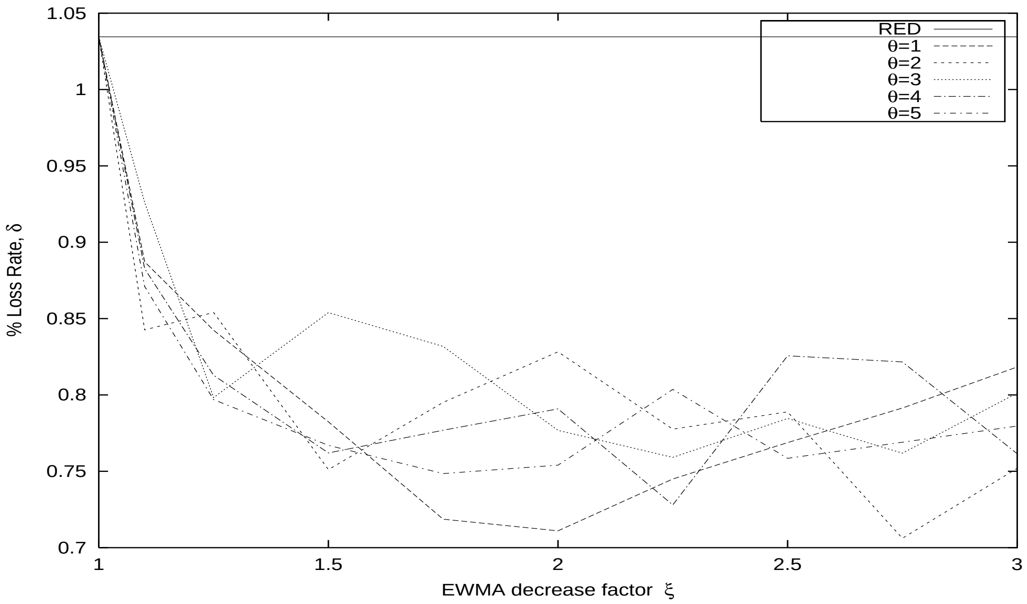

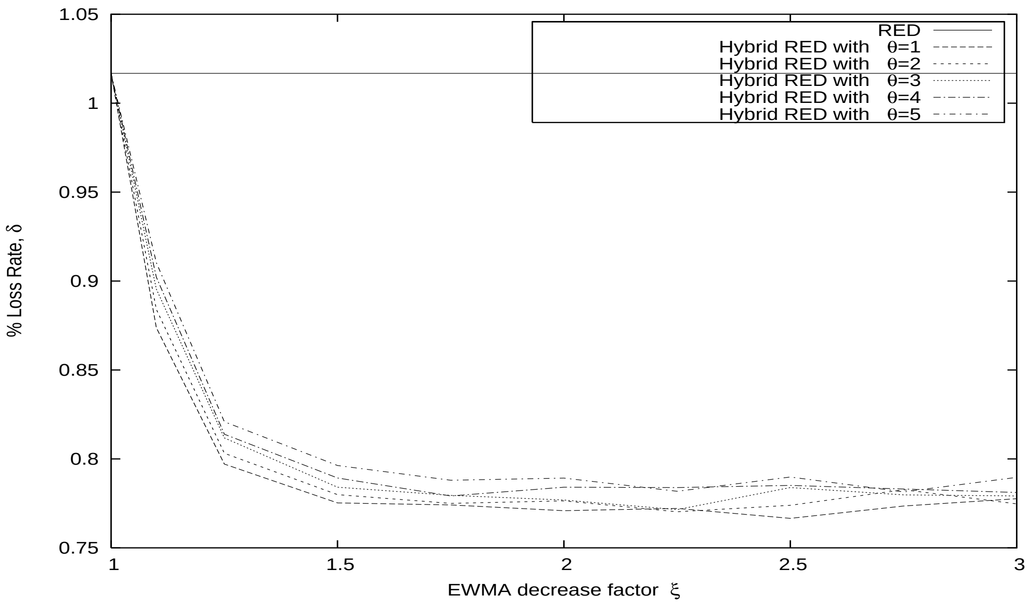

Analysis of TCP: results of one simulation vs means over 100 simulations

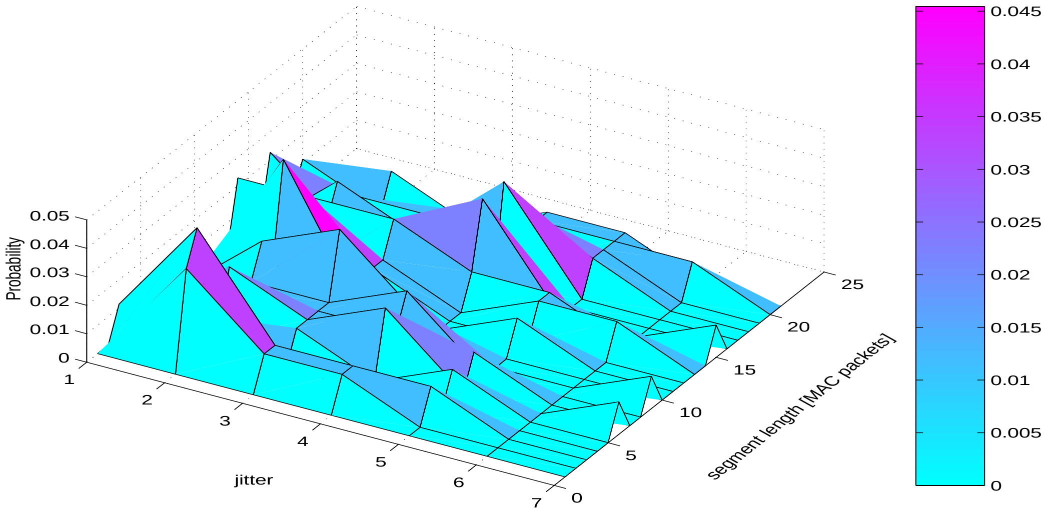

(a)

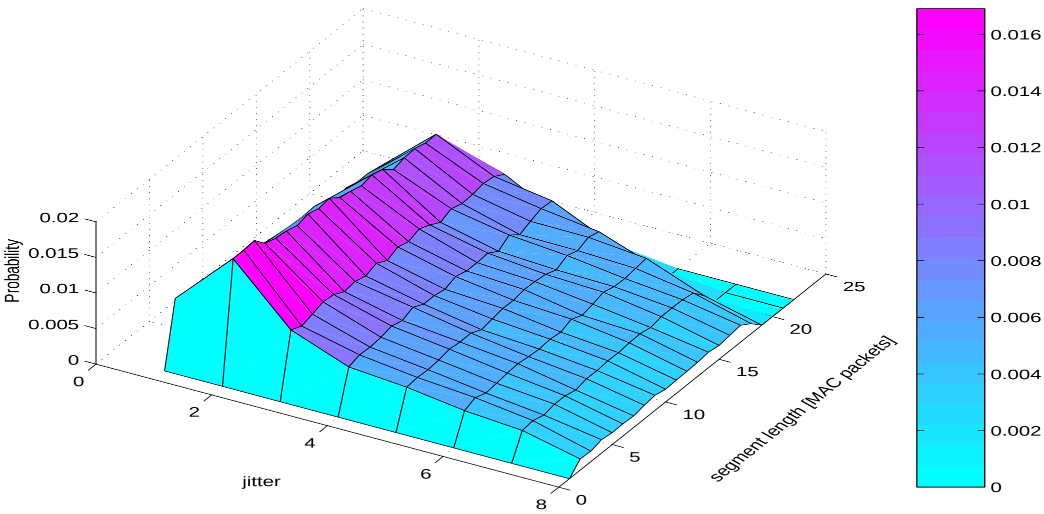

(b)

Evaluation of a Medium Access Control protocol of a mobile communication network. Results with relative statistical errors not greater than (a) 25%, (b) 1%, at 0.95 confidence lavel.

The only practical solution for producing credible results from simulation:

sequential on-line data analysis of statistical errors at consecutive checkpoints of simulation

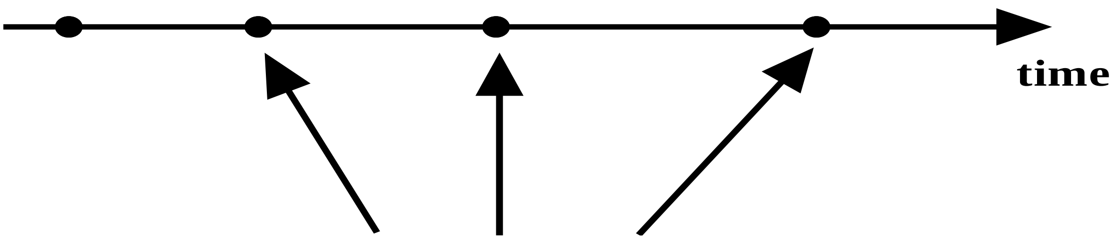

checkpoints
of output data analysis

the length of simulation is determined during simulation: the simulation is continued until statistical errors of results become satisfactorily small

Are simulation results reported in scientific literature appropriately analysed ?

Results of a survey
of 2246 publications on networks
published between 1992-1998 in:

(i) Proc. of IEEE INFOCOM
(ii) IEEE Trans. Communications
(iii) ACM/IEEE Trans. Networking
(iv) Performance Evaluation Journal

show that,
over 51% papers reported simulation results

see *[Pawlikowski et al., IEEE Comms., Jan. 2002.]*

Survey of 2246 papers on networks from

ACM/IEEE Trans. Networking	(1996-98, 223 papers)
IEEE Trans. Comms.	(1996-98, 657 papers)
Performance Evaluation J.	(1996-98, 174 papers)
Prc	

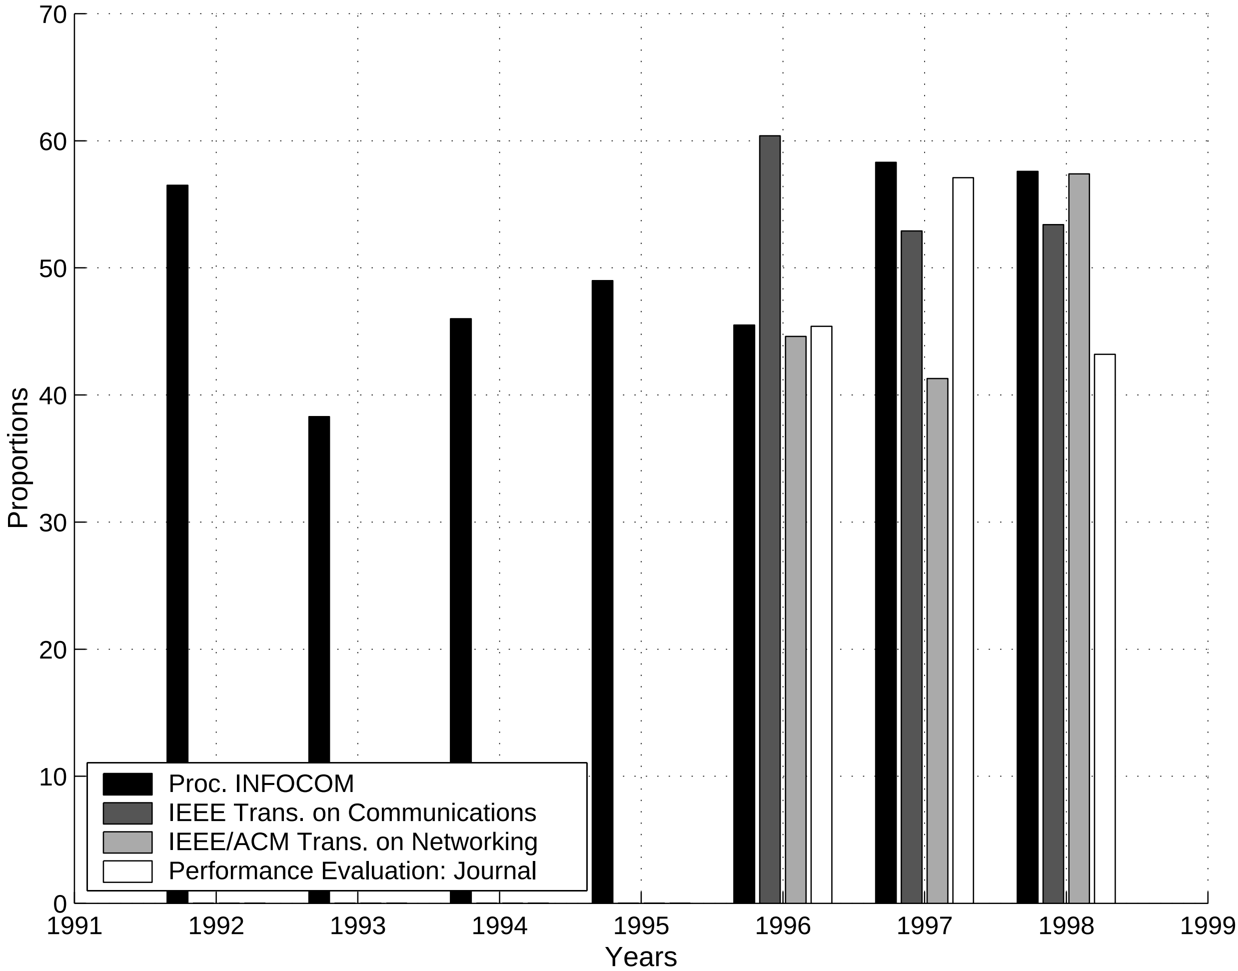

More than 51% of all papers reported results obtained by simulation

Pawlikowski et al. IEEE Comms., Jan. 2002

Proportion of papers reporting statistically analysed simulation results.

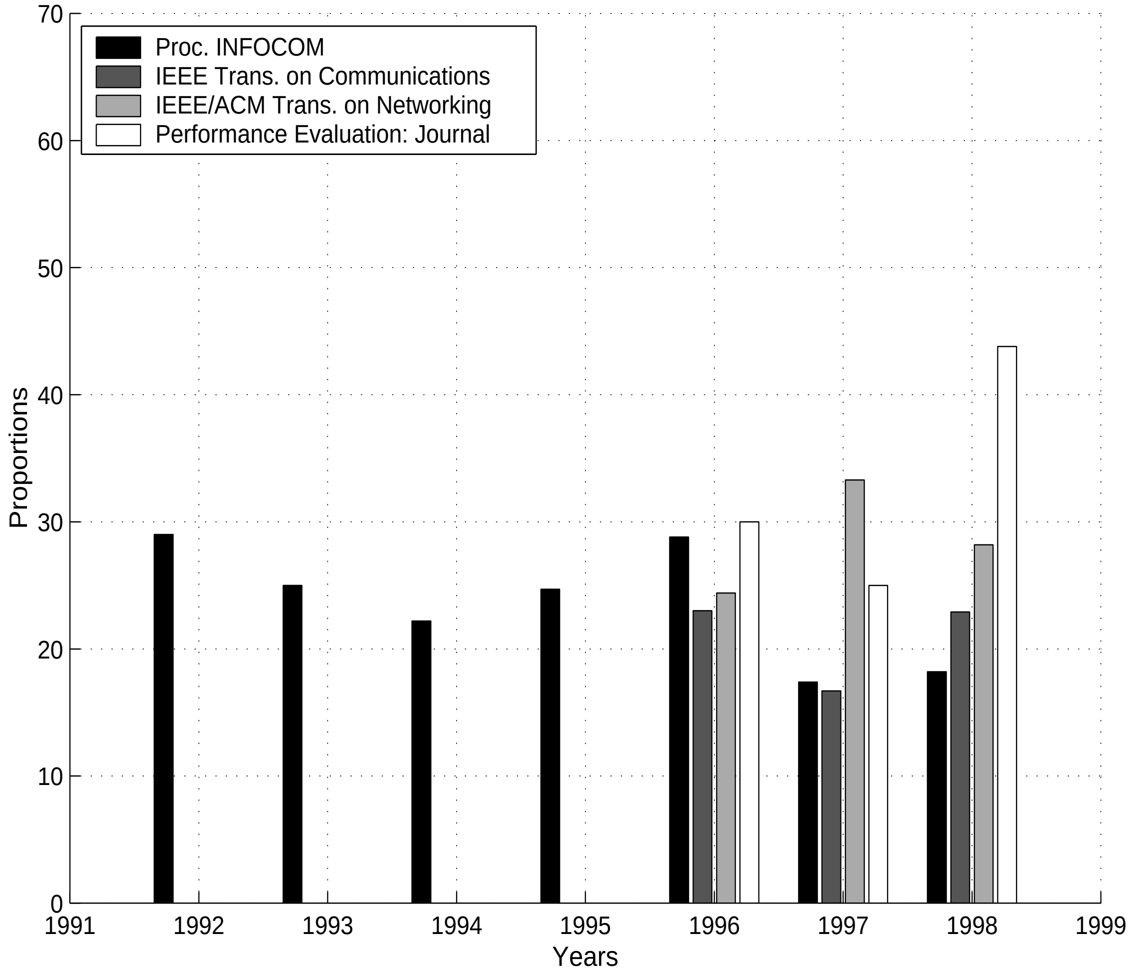

77% of papers that reported results based on simulation gave no evidence at all that simulation output data were statistically analysed.

Pawlikowski et al., IEEE Comms., Jan. 2002.

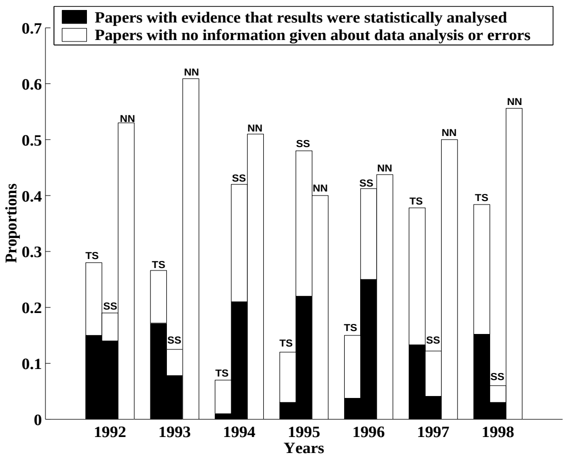

(a) Proc. of IEEE INFOCOM

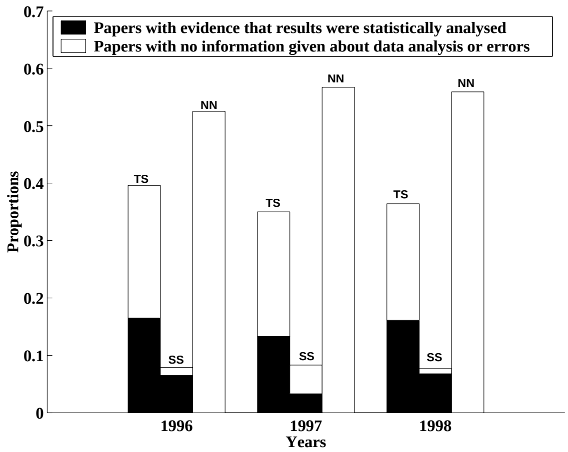

(b) IEEE Trans. Communications

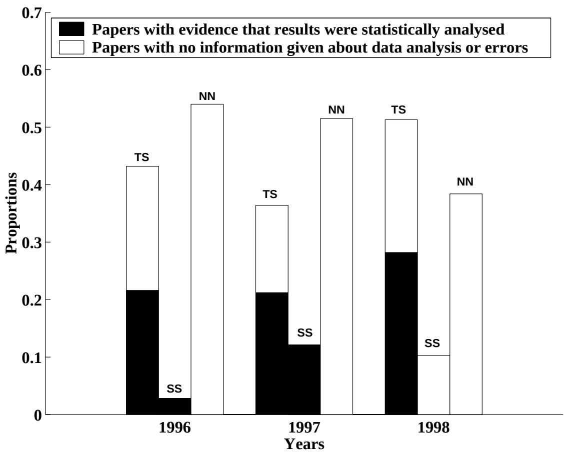

(c) ACM/IEEE Trans. Networking

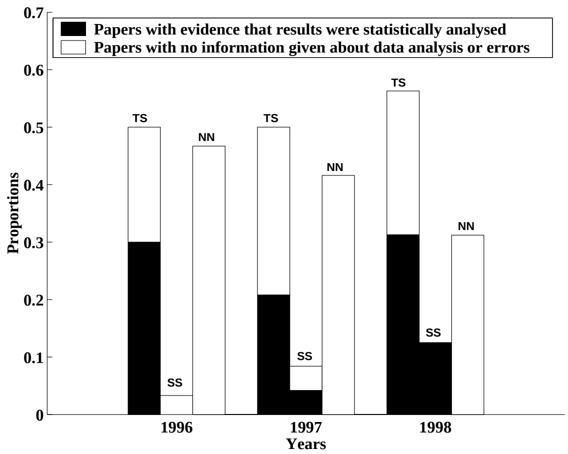

(d) Performance Evaluation Journal

Proportion of all papers on stochastic simulation that reported results analysed statistically:
from Terminating Simulation (**TS**), and
from Steady-State Simulation (**SS**).
NN = No information about the type of simulation nor about analysis of results

Results of this survey:

77% of papers that reported results based on simulation either

- gave no details of the simulation, or
- did not inform on how data were analysed (were they analysed at all ?), or
- both.

While all such papers reported non-repeatable simulations, some of them reported purely random results only !

Another survey:

Survey of all **151 papers published in Proceedings of the ACM Int. Symposium on Mobile Ad Hoc Networking and Computing** in 2000-2005:

- 75.5% of the papers reported results from simulation

- less than 15% of the papers reported repeatable studies of networks

- 12% of papers reported results which were based on sound statisitcal analysis of simulation output data.

See: S. Kurowski, T. Camp and M. Colagrosso.
"MANET Simulation Studies: The Incredibles".

Simulation studies of telecommunication networks:

- PSEUDO-SCIENCE ? (results based on personal beliefs of their authors)
- SCIENCE ? (results obtained and reported according to the principles of the scientific method)

... any new scientific hypothesis should be at least independently testable ...

Karl Popper (1901-1994), philosopher of science

Guidelines for presenting a credible publication based on simulation results:

- make sure that your simulation program represents valid simulation model;
- use well tested PRNG with sufficiently long cycle
- describe your statistical analysis of simulation output data in the paper, or refer to a technical report

Resolution to the credibility crisis of quantitative simulation:

- professional certification of simulation specialists ?
- automated sequential simulation ?

See *www.simprofessional.org* for discussion on professional certification of simulation specialists, proposed by Society for Computer Modelling and Simulation Int.

Problems of quantitative simulation studies of telecommunication networks:

- **limitations of statistics:** limited spectrum of methods for sequential analysis of simulation output data
- **technological limitations:** prohibitively long simulation time can be needed for securing sufficiently low statistical errors

For example, steady-state mean delay in $M/M/1/\infty$ buffer, with 5% error, at 0.95 confidence level, needs:

(a) **over 61 000 000** observations ($\rho = 0.99$),

(b) **over 6 143 000 000** observations

if ρ = 0.999.

Possible solution:

application of distributed processing (distributed generation of simulation output data)

Multiple Replications in Parallel (MRIP):

Multiple simulation engines independently produce statistical data by running their own simulation replications, but cooperate together in data analysis.

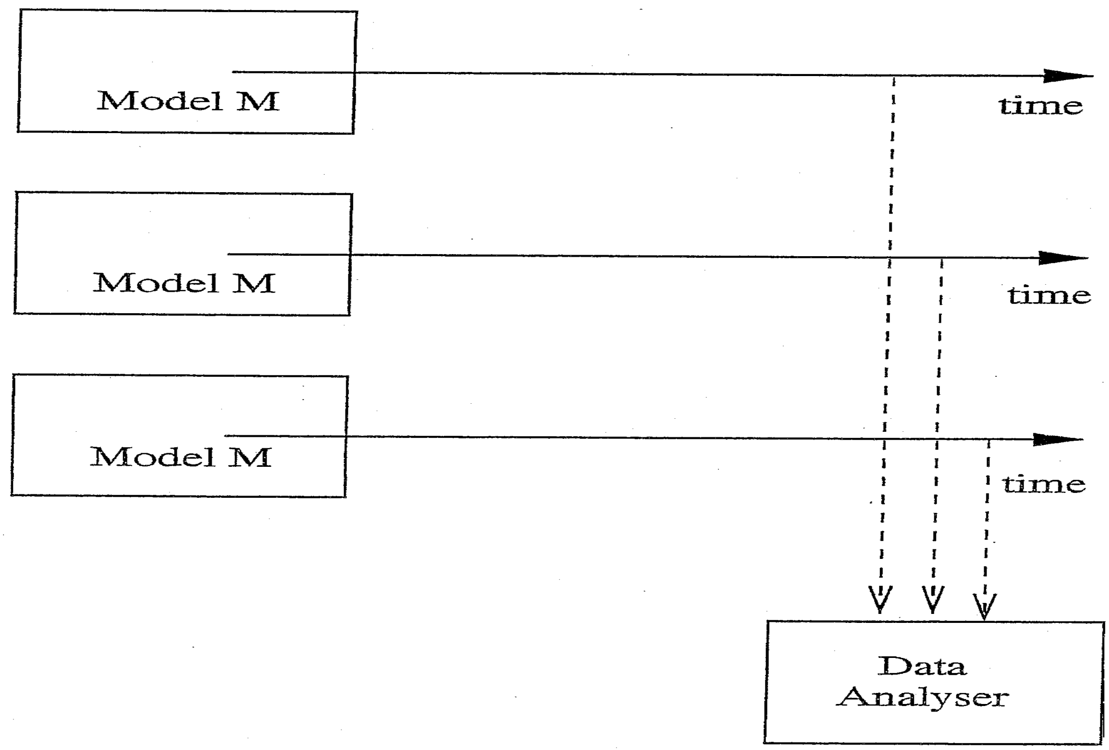

Main properties:

- model-independent parallelisation
- no parallel programming needed
- possible automation of parallelisation

MRIP as a (distributed) simulation in multiple time streams:

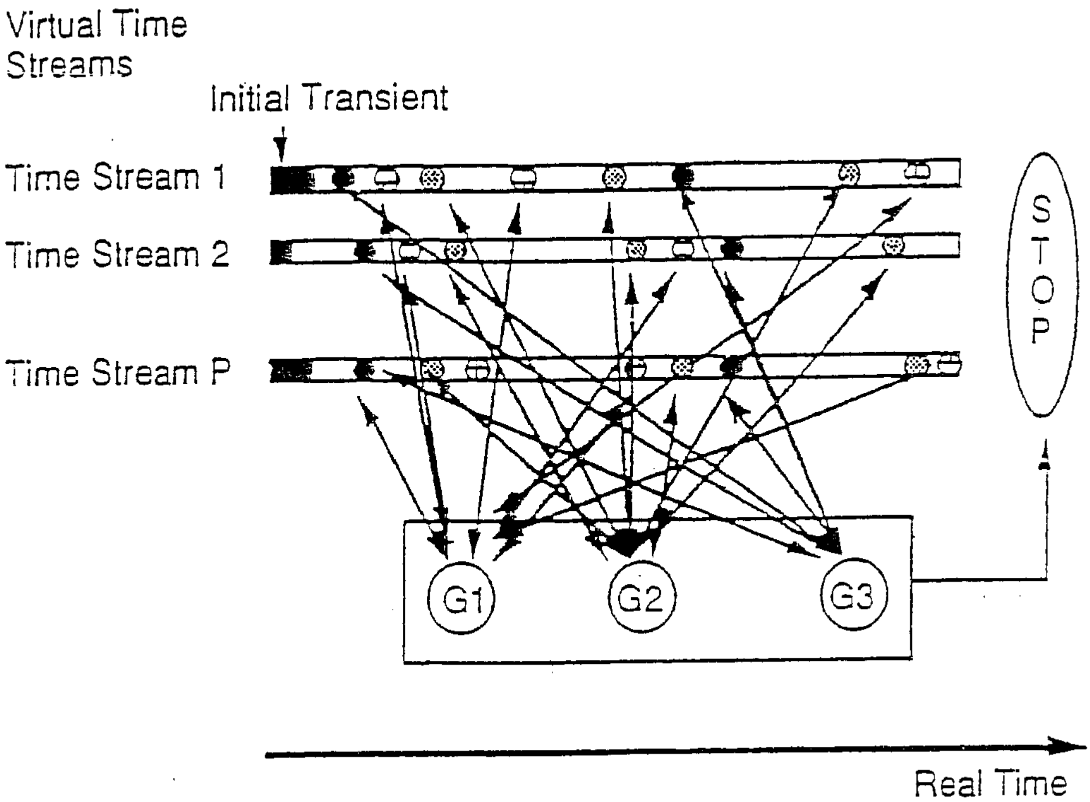

(⊖ = local checkpoint parameter 1,
⊛ = local checkpoint parameter 2,
● = local checkpoint parameter 3,
Gi = Global Estimation of parameter i, i=1,2,3.)

Truncated Amdahl's law of speedup in MRIP

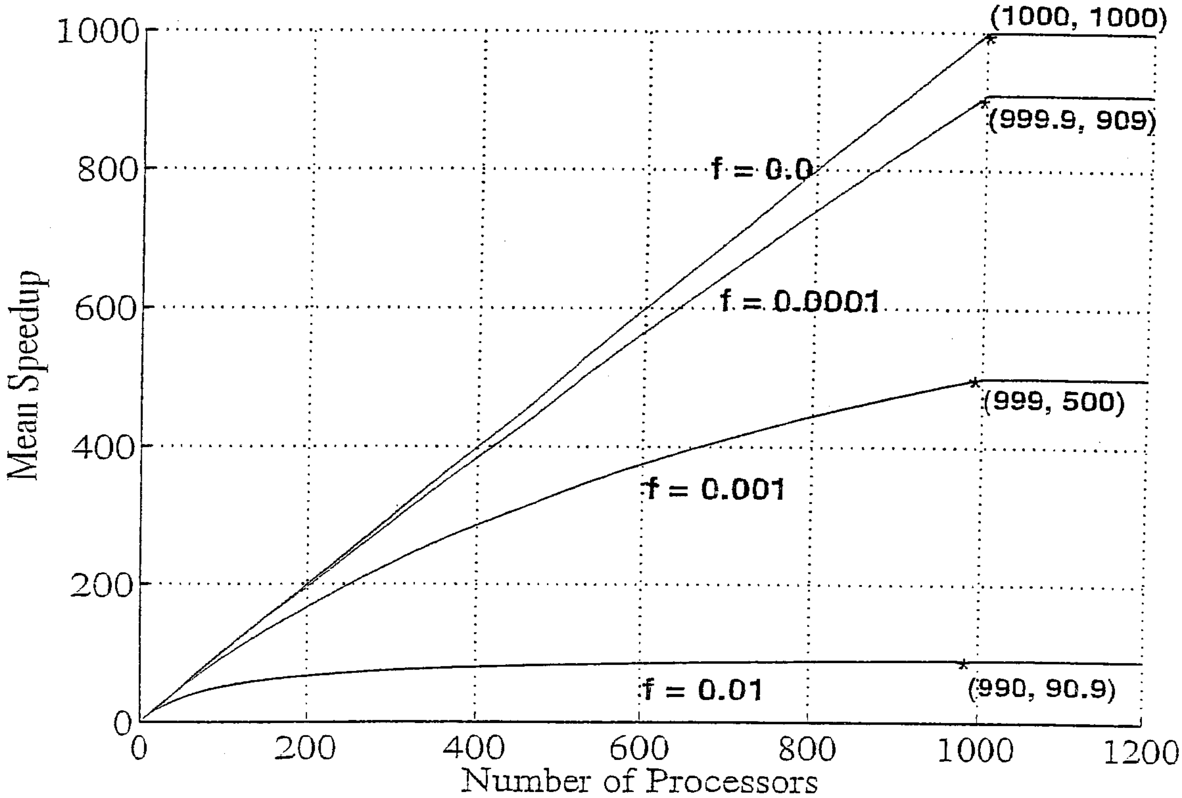

assumption: 1000 checkpoints needed
f = relative length of non-parallelisable stage of simulation (for example: the length of initial transient phase in steady-state simulation)

See *K. Pawlikowski and D. McNickle.*

"Speeding Up Stochastic Discrete-Event Simulation".

Proc. ESS'2001, Marseille, October 2001.

IMPLEMENTATION OF MRIP

AKAROA 2 = p**A**c**K**age for **A**utomatic gene**R**ation and process c**O**ntrol of p**A**rallel stochastic simulation (version **2**)

written in C++

tested on Sparc/SunOS4 & Solaris2, i386/Solaris2 and Linux

automatic parallelisation

automatic data collection and on-line analysis

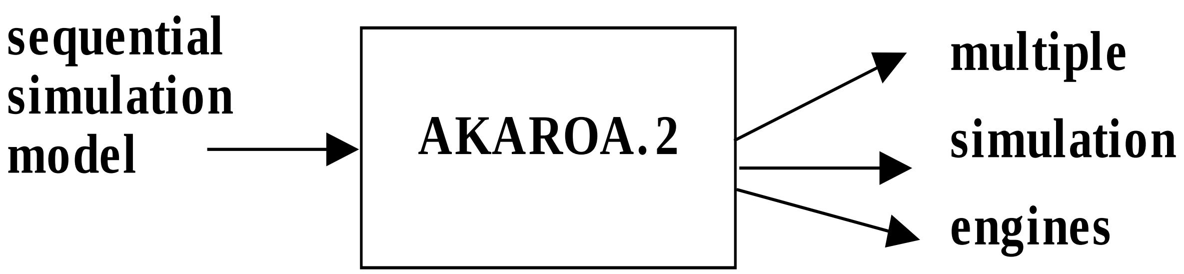

First reported automated implementations of MRIP:

EcliPSe [Rego & Sunderam, 1991]

AKAROA [Pawlikowski & Yau, 1992]

Example:

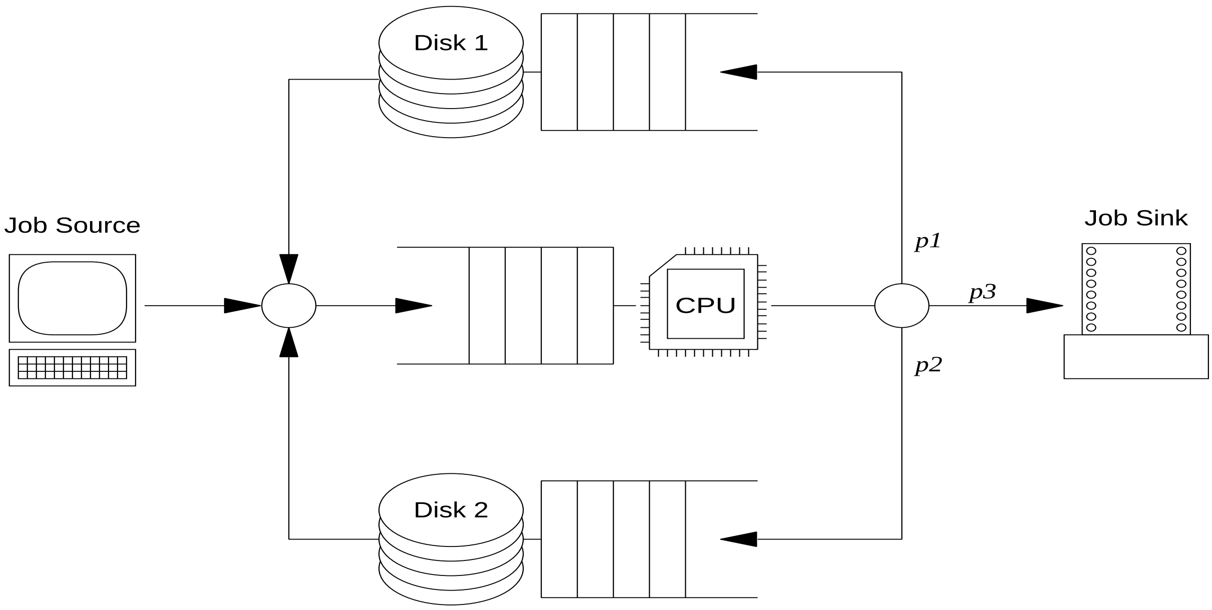

In the system loaded at 99%, analysis of the mean processing time of a job,
with error not exceeding 5%,
at 0.95 confidence level, requries collection of about 50 000 000 observations.

```
/*  *   mm1.C - M/M/1 Queueing System   */
#include "akaroa.H"
#include "akaroa/distributions.H"
#include "akaroa/process.H"
#include "akaroa/resource.H"

double arrival_rate; // Rate at which customers arrive
double service_rate; // Rate at which customers are served

// Server, modeled as a Resource with a capacity of 1 unit.

Resource server(1);

class Customer : public Process {
public:
  void LifeCycle();
};

void Customer::LifeCycle() {
  Time arrival_time, time_in_system;
  arrival_time = CurrentTime();
  server.Acquire(1);
  Hold(Exponential(1/service_rate));
  server.Release(1);
  time_in_system = CurrentTime() - arrival_time;
  AkObservation(time_in_system);
}
// The main program. After getting the load from the
// command line and calculating the arrival and service
// rates, we enter a loop generating new customers
// at the arrival rate.
int main(int argc, char *argv[]) {
  real load = atof(argv[1]);
  service_rate = 10.0;
  arrival_rate = load * service_rate;
  for (;;) {
    new Customer;
    Hold(Exponential(1/arrival_rate));
}}
```

Main functional blocks of AKAROA.2

akrun (shell command interface)

- **initiates simulation**
- **reports results to user**

akmaster (Parallel Simulation Manager)

- **manages launching of simulation engines**
- **receives checkpoint data and calculates global estimates**
- **controls simulation run length**

akgui: AKAROA.2 graphical user interface

Interprocess Communication
based on UNIX Internet-domain stream sockets

BLOCK DIAGRAM OF AKAROA.2

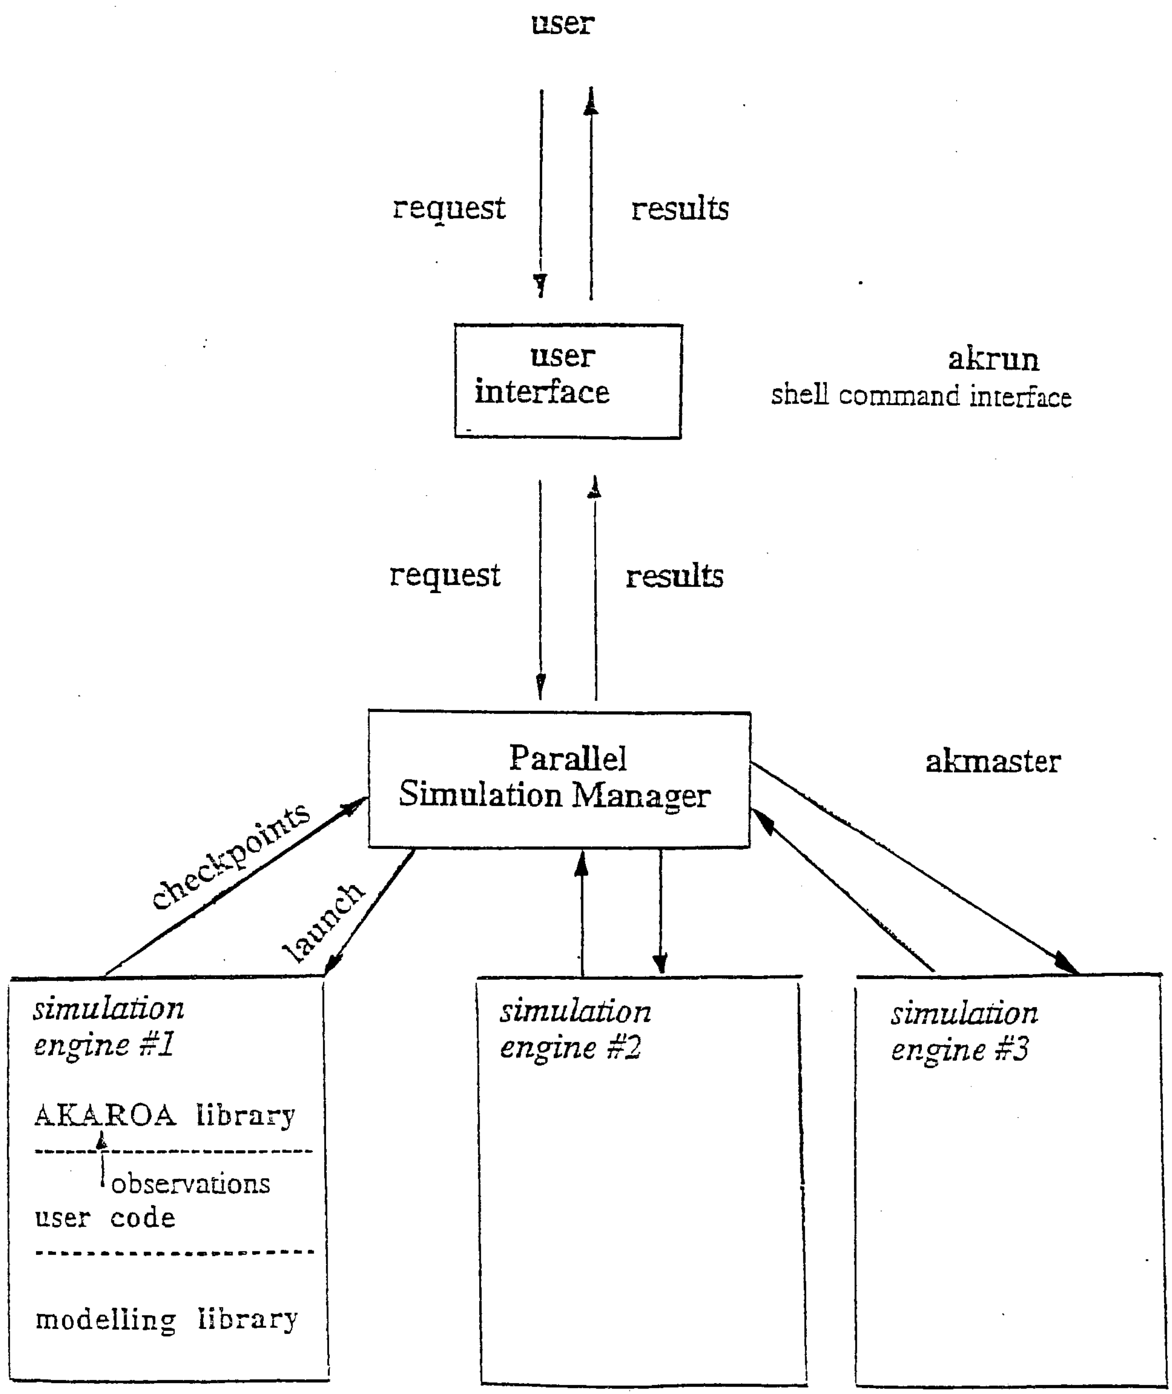

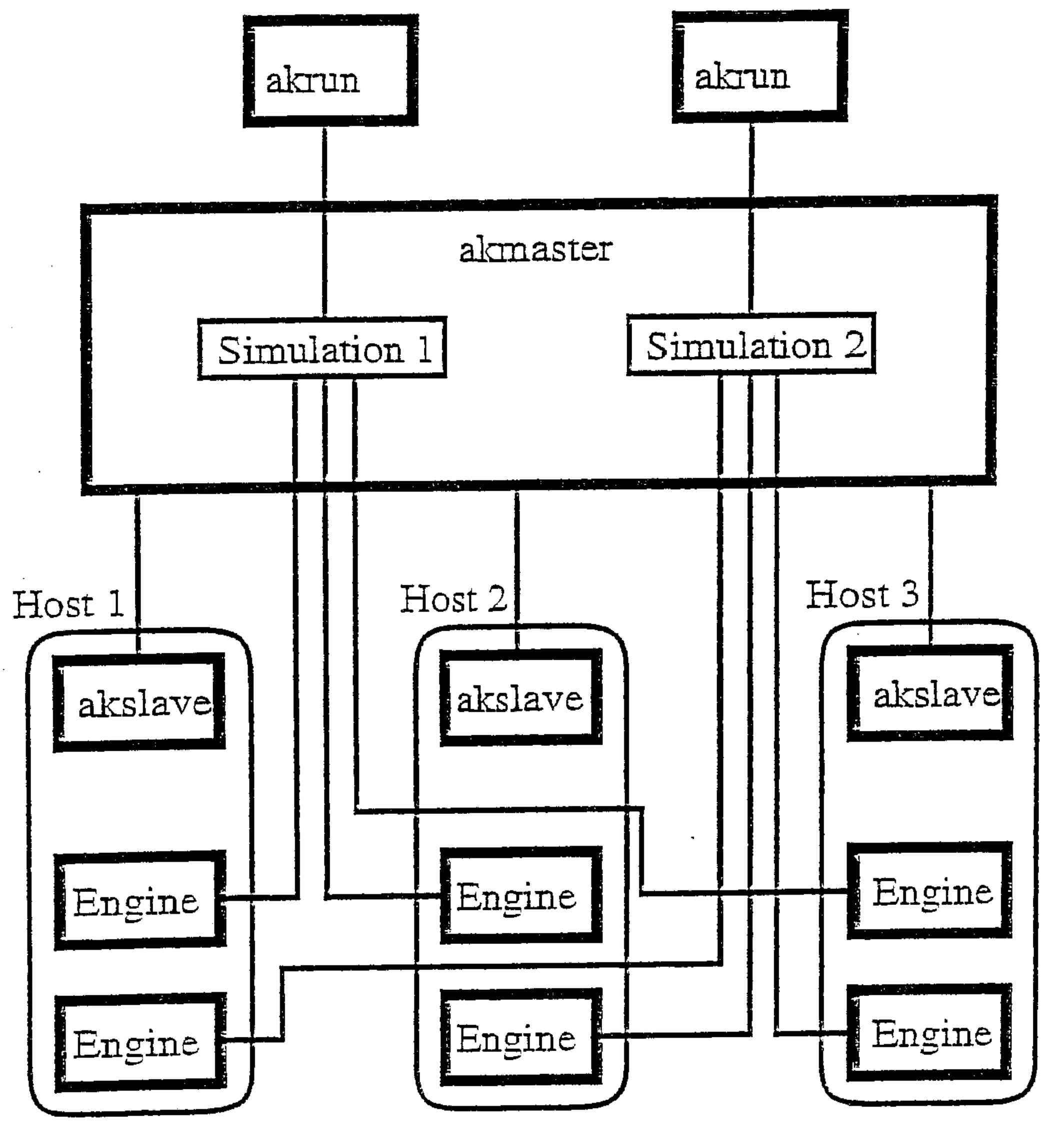

AKAROA.2 architecture

(2 simulations in progress, each one with 3 engines run on separate network hosts).

Akgui:

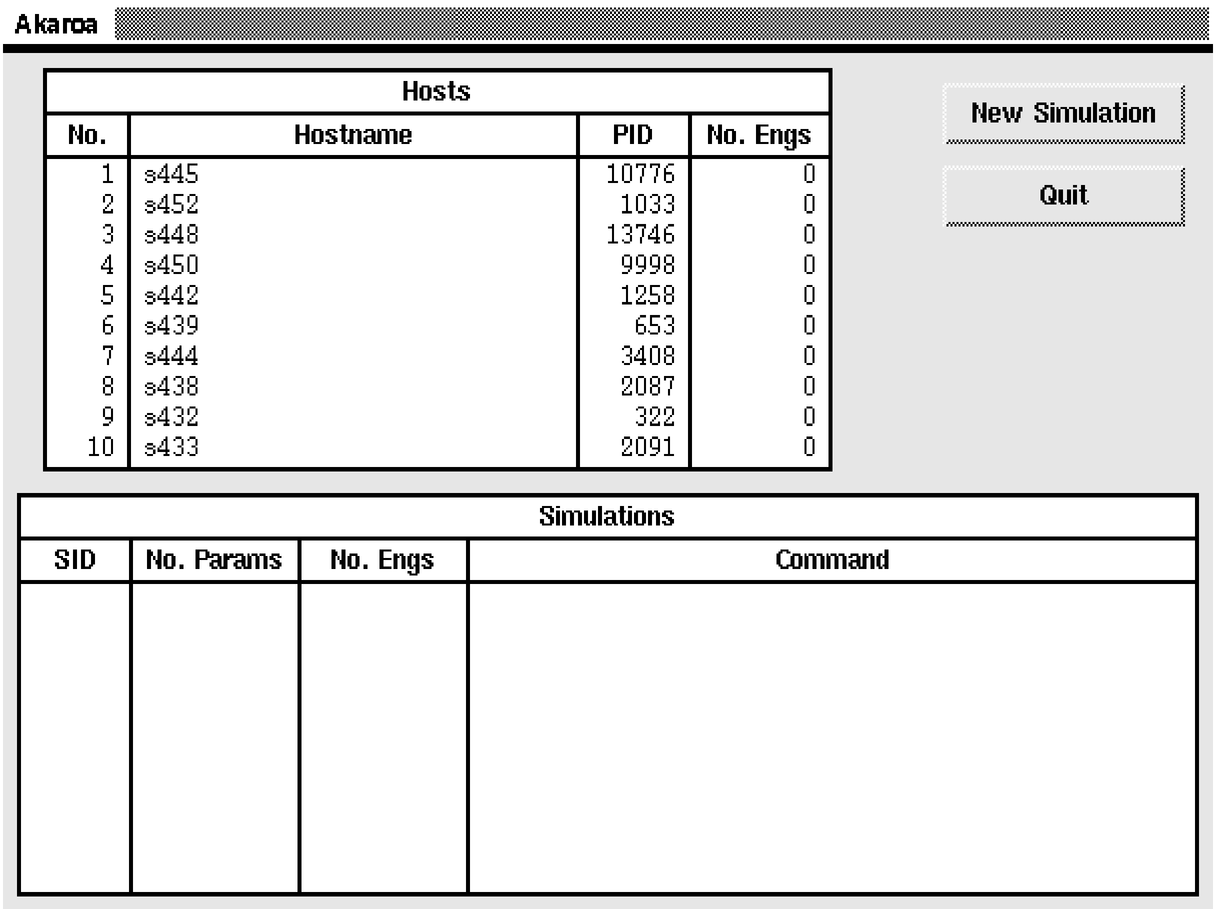

Hosts			
No.	Hostname	PID	No. Engs
1	s445	10776	0
2	s452	1033	0
3	s448	13746	0
4	s450	9998	0
5	s442	1258	0
6	s439	653	0
7	s444	3408	0
8	s438	2087	0
9	s432	322	0
10	s433	2091	0

Simulations			
SID	No. Params	No. Engs	Command

Akgui: before launching new simulation

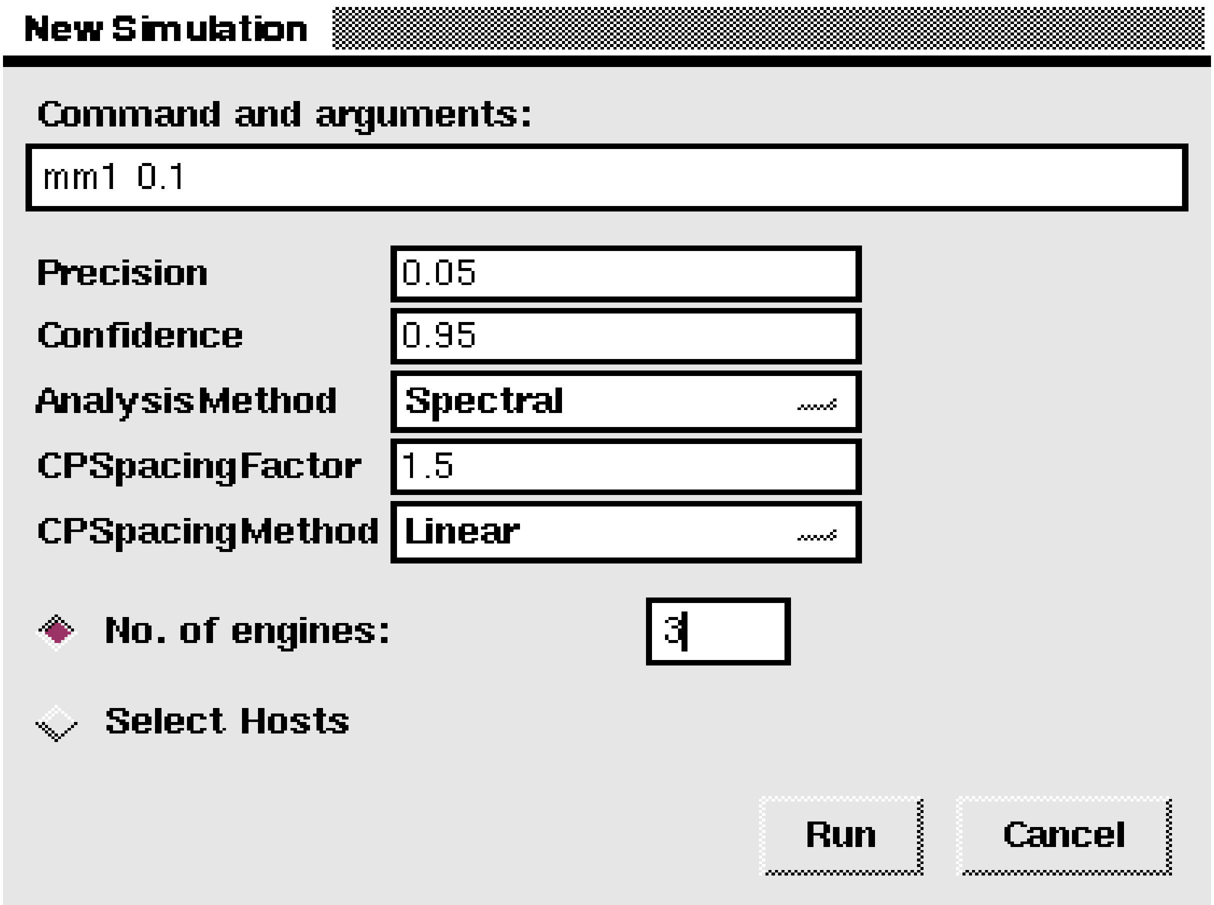

Akgui: during simulation

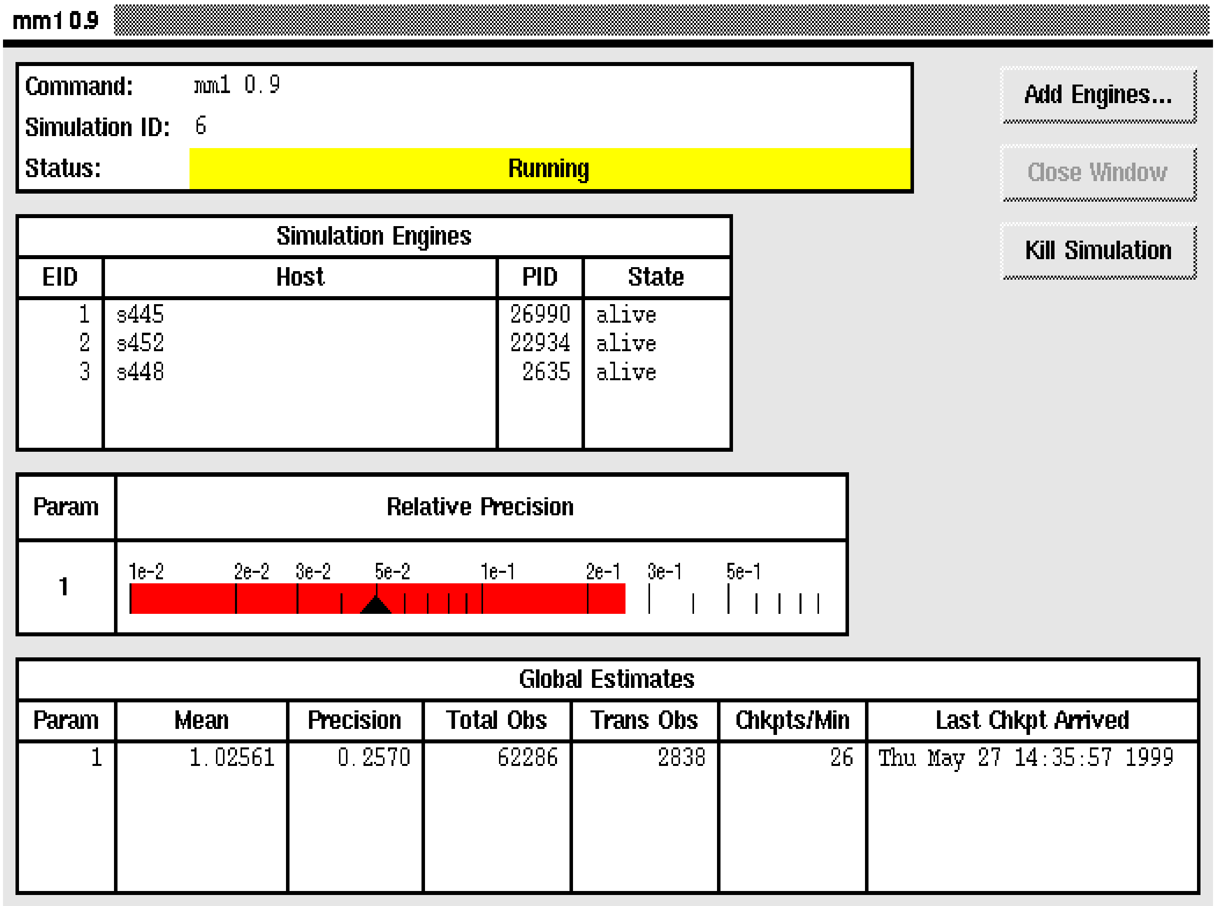

Akgui: the end of simulation

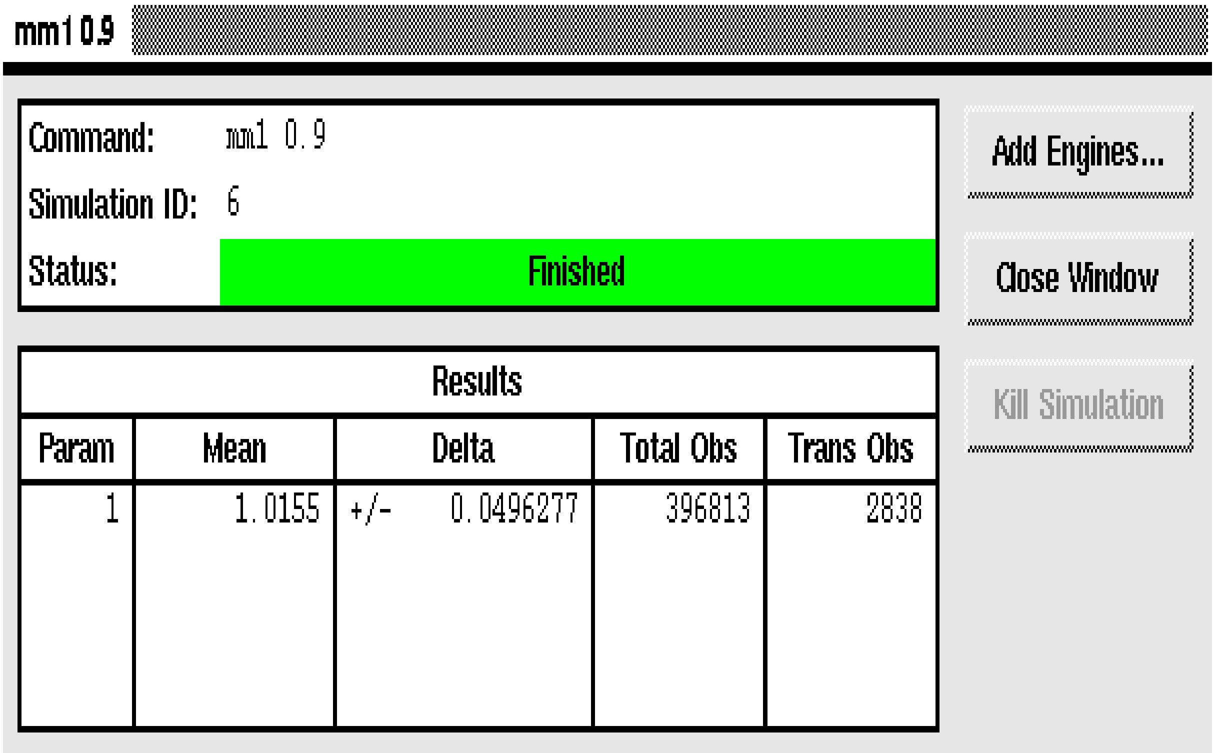

Some special features of AKAROA2:

- dynamic parallelisation
- fault tolerance (*if simulation engine is lost, simulation continues with remaining engines*)

AKAROA 2 as MRIP controller of other simulation packages

- AKAROA2/Ptolemy interface
  (1997, Ewing and Pawlikowski, University of Canterbury, New Zealand)

- AKAROA2/NS2 interface
  (2000, TU Berlin, Telecommunication Networks Group)

- AKAROA2/OMNET++ interface
  (2002, TU Berlin, Telecommunication Networks Group)

Akaroa 2 is offered as a free-ware
for non-profit research at universities.

See *www.cosc.canterbury.ac.nz/∼krys.*

A counter on Akaroa2 webpage has recently showed **about 11 000 visits** and **about 800 downloads of the code** (during last 5 years)

Akaroa2 is used at universities in over 60 countries (Jan. 2006), including:

Australia
Belarus
Belgium
Bosnia
Brazil
Canada
Chile
China
Colombia
Cyprus
Denmark
England
Finland
France
Germany
Greece
Guana
Hong Kong
Hungary
India
Indonesia
Iran
Ireland
Italy
Jakarta
Japan
Jordan
Korea
Lithuania
Macedonia
Malaysia
Mexico
New Zealand
Nigeria
Norway
Pakistan
Poland
Russia
Singapore
Slovakia
South Africa
Spain
Switzerland
Taiwan
Thailand
The Netherlands
Trinidad
Tunisia
Turkey
Ukraine
United Kingdom
United States
Vietnam

Commercial users include **IBM Research Labs in Zurich, Switzerland.**

Currently:
Akaroa2 can automatically produce accurate estimates of mean values, probabilities and proportions.

Work in progress:
sequential analysis of quantiles and probability distributions.

Future work:

- sequential analysis of rare events
- MRIP of distributed simulation models (distributed simulation on multiple clusters of workstations)
- distributed optimization in MRIP scenario.