ReSIST: Resilience for Survivability in IST

A European Network of Excellence

Contract Number: 026764

Deliverable D9: First Open Workshop report

Report Preparation Date: April 2007
Classification: Public

Contract Start Date: 1st January 2006
Contract Duration: 36 months
Project Co-ordinator: LAAS-CNRS
Partners: Budapest University of Technology and Economics
City University, London
Technische Universität Darmstadt
Deep Blue Srl
Institut Eurécom
France Telecom Recherche et Développement
IBM Research GmbH
Université de Rennes 1 – IRISA
Université de Toulouse III – IRIT
Vytautas Magnus University, Kaunas
Fundação da Faculdade de Ciencas da Universidade de Lisboa
University of Newcastle upon Tyne
Università di Pisa
QinetiQ Limited
Università degli studi di Roma "La Sapienza"
Universität Ulm
University of Southampton

Contents

1- Summary

The workshop was held at the Budapest University of Technology and Economics, on 21 and 22 March 2007.

The workshop was aimed at presenting salient results of the first year of activity, and to invite comments, criticisms, and suggestions for future investigation.

After a welcome address by Andras Pataricza (Budapest University of Technology and Economics), an overview of ReSIST by Jean-Claude Laprie (LAAS-CNRS) presented the network objectives and the progresses made so far.

Presentations by ReSIST members include a selection of topics from the State of Knowledge document produced by the Network, and the demonstration of an ontology-based resilience knowledge base. The presentation titles are as follows:

- Data distribution in large-scale systems, by Roberto Baldoni (Università degli studi di Roma "La Sapienza")
- Cooperative backup in dynamic systems, by Marc-Olivier Killijian (LAAS-CNRS)
- Challenges and advances in dependable e-voting systems: technical and socio-technical aspects, by Peter Ryan (University of Newcastle upon Tyne) and Lorenzo Strigini (City University, London)
- Modeling and evaluation of largeness in evolving systems, by Andrea Bondavalli (Università di Firenze)
- Towards attack modelization thanks to honeypot data processing, by Marc Dacier (Institut Eurécom)
- Scalable verification of systems with cryptography, by Birgit Pfitzmann (IBM Research Zurich)
- Prototype knowledge base: an on-line information service in dependability and security, by Hugh Glaser (University of Southampton)

It has to be noted that the six presentations selected from the State of Knowledge document resulted from a rather drastic selection, as the document is composed of 22 chapters covering the design, the verification, and the evaluation of resilient computer systems.

Those presentations were complemented by

- two invited talks:
  - Probabilistic Validation of Computer System Survivability, by Bill Sanders (University of Illinois at Urbana-Champain)
  - Modelling of failures: from chains to coincidences, by Erik Hollnagel (Ecole des Mines, Sophia Antipolis)

and by

- a panel moderated by Luca Simoncini (University of Pisa) where selected European projects presented their views of resilience:
  - DESEREC Integrated Project, Benoît Bruyère (Thales),
  - ESFORS Coordination Action, Aljosa Pasic (Atos Origin),
  - SERENITY Integrated Project, Domenico Presenza (Ingegneria Informatica),
  - HIDENETS Specific Targeted Research Project, Hans Peter Schwefel (Aalborg University).

The concluding session, moderated by Tom Anderson (University of Newcastle upon Tyne), was an opportunity for the attendees to give their viewpoints.

The workshop was attended by 93 persons:

- 73 members of ReSIST,
- the project officer and the 3 reviewers,
- 5 members of the scientific council, one of them being an invired speaker,
- the other invited speaker,
- the 4 panelists,
- 6 additional external attendees.

The remainder of this report gives:

1) The workshop programme.
2) The attendance list.
3) The copies of the slides presented during the workshop.

2- Programme

ReSIST: Resilience for Survivability in IST

A European Network of Excellence

http://www.resist-noe.org

First Open Workshop

Budapest University of Technology and Economics

21-22 March 2007

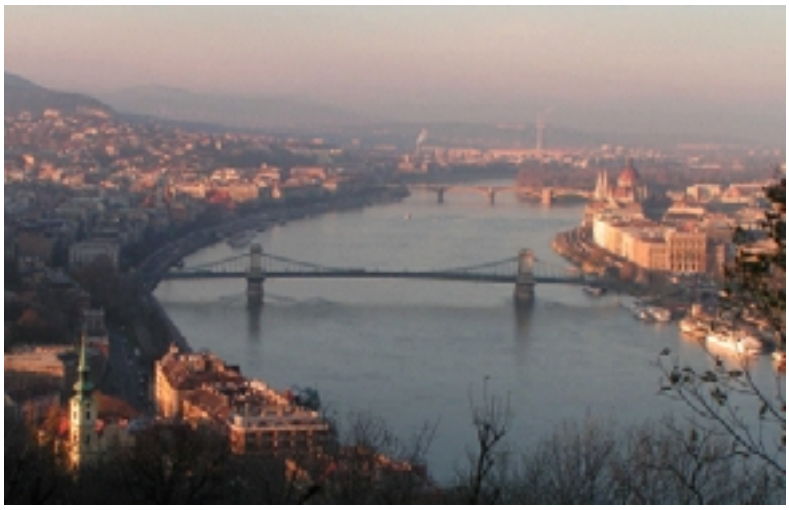

Network Partners

- LAAS-CNRS, Toulouse, France (Coordinator)
- Budapest University of Technology and Economics, Hungary
- City University, London, UK
- Technische Universität Darmstadt, Germany
- Deep Blue Srl, Roma, Italy
- Institut Eurécom, Sophia Antipolis, France
- France Telecom Recherche et Développement, Lannion and Caen, France
- IBM Research Gmbh, Zürich, Switzerland
- Université de Rennes 1 – IRISA, France
- Université Paul Sabatier Toulouse III – IRIT, France
- Vytautas Magnus University, Kaunas, Lithuania
- Fundação da Faculdade de Ciencas da Universidade de Lisboa, Portugal
- University of Newcastle upon Tyne, UK
- Università di Pisa, Italy
- QinetiQ Limited, Malvern, UK
- Università degli studi di Roma "La Sapienza", Italy
- Universität Ulm, Germany
- University of Southampton, UK

Information Society Technologies

SIXTH FRAMEWORK PROGRAMME

Contract Number: 026764

About ReSIST

ReSIST is a Network of Excellence that addresses the strategic objective “Towards a global dependability and security framework” of the European Union Work Programme, and responds to the stated “need for resilience, self-healing, dynamic content and volatile environments”.

It integrates leading researchers active in the multidisciplinary domains of Dependability, Security, and Human Factors, in order that Europe will have a well-focused coherent set of research activities aimed at ensuring that future “ubiquitous computing systems” – the immense systems of ever-evolving networks of computers and mobile devices which are needed to support and provide Ambient Intelligence (AmI) – have the necessary resilience and survivability, despite any physical and residual development faults, interaction mistakes, or malicious attacks and disruptions.

About the Workshop

ReSIST started on January 2006. The workshop is aimed at presenting salient results of the first year of activity, and to invite comments, criticisms, and suggestions for future investigation.

Presentations by ReSIST members include a selection of topics from the State of Knowledge document produced by the Network, and the demonstration of an ontology-based resilience knowledge base.

Those presentations are complemented by

- two invited talks by distinguished and highly renowned speakers,

and by

- a panel where selected European projects will present their views of resilience, and compare them to ReSIST's views.

Programme

Wednesday 21 March

12h Registration

12h30 - 14h Lunch

14h - 14h35 ***Opening Session***
Session Chair: Andras Pataricza (Budapest University of Technology and Economics)
ReSIST: resilience for survivabilty, an overview, Jean-Claude Laprie (LAAS-CNRS)

14h35 - 16h05 ***Resilience Design***
Session Chair: Michel Raynal (Université de Rennes I - IRISA)
Data distribution in large-scale systems, Roberto Baldoni (Università degli dtudi di Roma "La Sapienza")
Cooperative backup in dynamic systems, Marc-Olivier Killijian (LAAS-CNRS)
Challenges and advances in dependable e-voting systems: technical and socio-technical aspects, Peter Ryan (University of Newcastle upon Tyne) and Lorenzo Strigini (City University, London)

16h05 - 16h35 Coffee Break

16h35 - 17h15 ***Invited Talk 1***
Session Chair: Algirdas Avizienis (Vytautas Magnus University, Kaunas)
Probabilistic Validation of Computer System Survivability, Bill Sanders (University of Illinois at Urbana-Champain)

17h15 - 18h30 ***Resilience Evaluation and Verification***
Session Chair: Karama Kanoun (LAAS-CNRS)
Modeling and evaluation of largeness in evolving systems, Andrea Bondavalli (Università di Firenze)
Towards attack modelization thanks to honeypot data processing, Marc Dacier (Institut Eurécom)
Scalable verification of systems with cryptography, Birgit Pfitzmann (IBM Research Zurich)

20h Workshop Banquet

Thursday 22 March

8h30 - 9h10 ***Resilience Knowledge Base***
Session Chair: Brian Randell (University of Newcastle upon Tyne)
Prototype knowledge base: an on-line information service in dependability and security, Hugh Glaser (University of Southampton)

9h10 - 9h50 ***Invited Talk 2***
Session Chair: Alberto Pasquini (Deep Blue)
Modelling of failures: from chains to coincidences, Erik Hollnagel (Ecole des Mines, Sophia Antipolis)

9h50 - 10h20 Coffee Break

10h20 - 11h20 ***Panel***
Resilience Views from other European Projects
Panel Moderator: Luca Simoncini (Unversità di Pisa)
Panelists:
Benoît Bruyère (Thales), *DESEREC Integrated Project*
Aljosa Pasic (Atos Origin), *ESFORS Coordination Action*
Domenico Presenza (Ingegneria Informatica), *SERENITY Integrated Project*
Hans Peter Schwefel (Aalborg University), *HIDENETS Specific Targeted Research Project*

11h20 - 12h30 ***Conclusions***
Session Chair: Tom Anderson (University of Newcastle upon Tyne)
Future research directions, structuring effect of ReSIST, Jean-Claude Laprie (LAAS-CNRS)
General discussion

12h30 - 14h Lunch

Registration fee

Registration fee is 100 Euros, covering

- attendance to the workshop,
- a CD containing the State of Knowledge on Resilience, produced by ReSIST,
- the Banquet, Lunches, Coffee Breaks.

No registration fee is required from students.

Attendance is limited. Registrations will be processed on a first-come first-served basis.

Fellowships

A limited number of fellowships will be made available for scientists and industrial experts from the New Member States.

Please apply in e-mail to:
resistmeeting@mit.bme.hu

Location

Budapest University of Technology and Economics

The ReSIST Workshop will take place in Builiding **A**

How to get to the University:
http://portal.bme.hu/langs/en/where.aspx

Map of the University: http://portal.bme.hu/terkep.aspx

Hotels

Danubius Hotel Gellért****

1111 Budapest, Szent Gellért tér 1.

Room reservation:
Phone: +36 1 889-5501
Fax: +36 1 889-5505
E-mail: gellert.reservation@danubiusgroup.com
http://www.danubiushotels.com/gellert

Hotel Mercure Duna***

1095 Budapest, Soroksári út 12.

Room reservation:
Phone: +36 1 455-8300
Fax: +36 1 455-8385
http://www.accorhotels.com/accorhotels/fichehotel/gb/mer/2025/fiche_hotel.shtml

BME Professor's Guesthouse***

1111 Budapest, Stoczek utca 5-7, 7th floor

Room reservation:
Phone.: +36 1 463-4103
Fax: +36 1 463-3936
http://www.otevszak.hu/hotel/angol/indexed.php

ReSIST Open Workshop

Budapest University of
Technology and Economics

Registration Form

Fax to +36 1 463 26 67 or email to resistmeeting@mit.bme.hu, before 5 March 2007

Attendee:

Name (First Last): _____

Email: _____

Company/Institution: _____

Address: _____

_____

City: _____ State/Province: _____

Country: _____ Zip/Postal Code: _____

Phone: _____ Fax: _____

Special Dietary Needs: _____

Registration fee: 100 EUR, covering

- attendance to the Workshop,
- a CD containing the State of Knowlegde on Resilience produced by ReSIST,
- the Banquet, Lunches, Coffee Breaks.

Students

No registration fee is required from students. If you are a student, please tick ❑
Evidence of student status will be requested upon registration.

Fellowships

A limited number of fellowships will be made available for scientists and industrial experts from the New Member States. *Please apply in e-mail to*: resistmeeting@mit.bme.hu

Payment

By Credit Card:

Card type: VISA ❑ EUROCARD/MASTERCARD ❑ DINERS CLUB ❑

Name of card holder: _____

Card number: _____ Expiry date: _____

CVV number (last three digits number at the back of the card): _____

By bank transfer:

IBAN: HU55 1091 8001 0000 0003 3926 0098
Name of the Bank: HVB Bank
Address of the Bank: H-1111 Lágymányosi u. 2
Swift code: BACXHUHB

Budapest University of Technology and Economics
H-1117 Budapest, Magyar tudósok krt. 2., Hungary
Phone: +36 1 463 35 82 Fax: +36 1 463 26 67
resistmeeting@mit.bme.hu

3- Attendance List

Last name	First name	Organisation
Alberdi	Eugenio	City University
Anderson	Tom	Newcastle University
Andrews	Zoe	Newcastle University
Avizienis	Algirdas	VMU, Kaunas, Lithuania
Bacivarov	Ioan C.	University "Politechnica" Bucharest
Baldoni	Roberto	University of Roma "La Sapienza"
Banâtre	Michel	IRISA-Rennes
Basnyat	Sandra	IRIT, Université Paul Sabatier
Benato	Roberto	University of Roma "La Sapienza"
Beraldi	Roberto	University of Roma "La Sapienza"
Bernardeschi	Cinzia	University of Pisa
Bokor	Peter	BUTE
Bondavalli	Andrea	University of Firenze
Bonomi	Silvia	University of Roma "La Sapienza"
Bruyere	Benoit	THALES
Bryans	Jeremy	Newcastle University
Carvalho	Pedro	University of Lisboa
Correia	Miguel	University of Lisboa
Courtès	Ludovic	LAAS-CNRS
Crouzet	Yves	LAAS-CNRS
Culo	Oliver	VMU, Kaunas, Lithuania
Dacier	Marc	Institute Eurecom
Debar	Hervé	France Telecom
Di Marzo Serugendo	Giovanna	Birkbeck College, UK
Faconti	Giorgio	University of Pisa
Fitzgerald	John	Newcastle University
Glaser	Hugh	University of Southampton
Gönczy	László	BUTE
Grigonyte	Gintare	VMU, Kaunas, Lithuania
Harrison	Michael	Newcastle University
Hollnagel	Erik	Pole Cindyniques
Horváth	Ákos	BUTE
Huszerl	Gábor	BUTE
Kaaniche	Mohamed	LAAS-CNRS
Kanoun	Karama	LAAS-CNRS
Killijian	Marc-Olivier	LAAS-CNRS
Knight	John	University of Virginia
Kocsis	Imre,	BUTE
Kovács	Máté	BUTE
Kövi	András	BUTE
Kurth	Helmut	Atsec
Lac	Chidung	France Telecom
Laprie	Jean-Claude	LAAS-CNRS
Laszlo Pasztor	Peter	BUTE
Leita	Corrado	Institut Eurecom
Long	Derek M.	CISA Ltd.
Majuntke	Matthias	TU Darmstadt
Majzik	István	BUTE
Martini	Luca	University of Pisa
Masci	Paolo	University of Pisa
Micskei	Zoltán	BUTE
Millard	Ian	University of Southampton
Moffat	Nick	QinetiQ
Morganti	Michele	Siemens
O'Halloran	Colin	QinetiQ
Paindaveine	Yves	European Commission
Palanque	Philippe	IRIT, Université Paul Sabatier
Pasic	Aljosa	Atos Origin

Pasquini	Alberto	Deep Blue
Pataricza	András	BUTE
Pfeifer	Holger	University of Ulm
Pfitzmann	Birgit	IBM Zurich Research Lab
Pinter	Gergely	BUTE
Popov	Peter	City University
Posegga	Joachim	University of Hamburg
Presenza	Domenico	Engineering
Ramanathan	Sakkaravarthi	France Telecom
Raynal	Michel	IRISA-Rennes
Riordan	James	IBM Zurich Research Lab
Roudier	Yves	Institute Eurecom
Roy	Matthieu	LAAS-CNRS
Rushby	John	Computer Science Laboratory
Ryan	Peter Y. A.	Newcastle University
Sanders	William	University of Illinois at Urbana-Champaign
Schiper	Andre	EPFL
Schöller	Markus	Lancaster University
Schwefel	Hans-Peter	Aalborg University
Scipioni	Sirio	University of Roma "La Sapienza"
Sidlauskas	Kestutis	VMU, Kaunas, Lithuania
Simoncini	Luca	University of Pisa
Stankovic	Vladimir	City University
Sterbenz	James	Lancaster University
Strigini	Lorenzo	City University
Stroud	Robert	Newcastle University
Suri	Neeraj	TU Darmstadt
Thomas	Martyn	Thomas Associates
Tirtea	Rodica	University of Oradea, Romania
Tóth	Dániel	BUTE
Urvoy-Keller	Guillaume	Institute Eurecom
van Moorsel	Aad	Newcastle University
Verissimo	Paulo	University of Lisboa
von Henke	Friedrich W.	University of Ulm
Waeselynck	Helene	LAAS-CNRS

4- Slides

ReSIST

Resilience for Survivability in IST

A European Network of Excellence

Information Society Technologies

SIXTH FRAMEWORK PROGRAMME

- Rationale
- Joint Programme of Activities and Logic
- Partnership
- Organisation
- First Year Results
- Open Workshop and Review
- About Resilience

Rationale

(Reasonably) known: High dependability
for safety-critical or availability-critical systems

Avionics, railway signalling, nuclear control, etc.

Transaction processing, back-end servers, etc.

Continuous complexity growth
Large, networked, evolving, applications running on open systems, fixed or mobile,
i.e., *ubiquitous systems*

Dependability gap between necessary trust for realistic AmI and operational statistics

Scalability of Dependability

In addition to rigorous functional design, provision of

Resilience for Survivability

Development or physical accidental faults

Malicious attacks

Interaction mistakes

Vulnerabilities

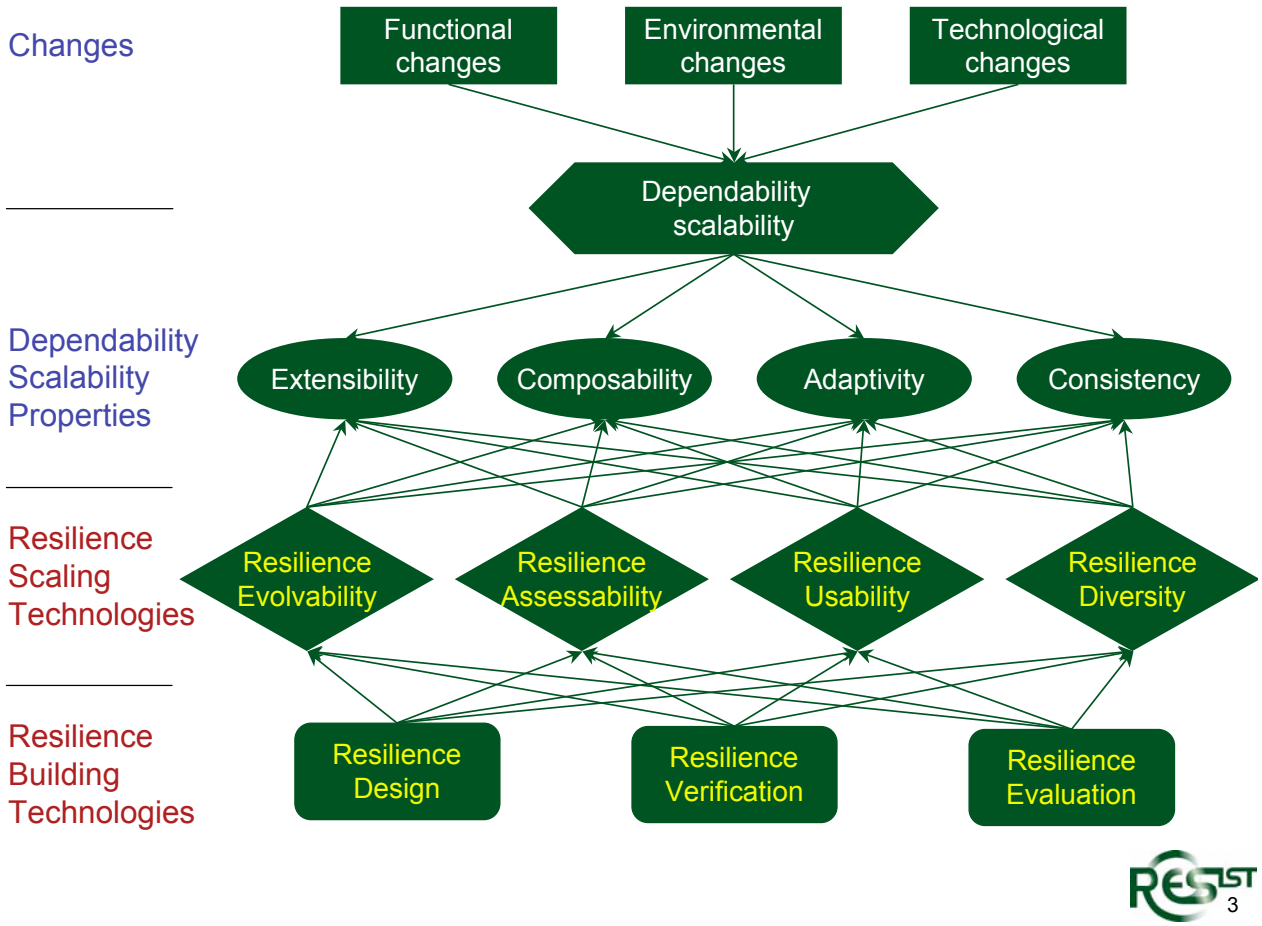

Joint Programme of Activities and Logic

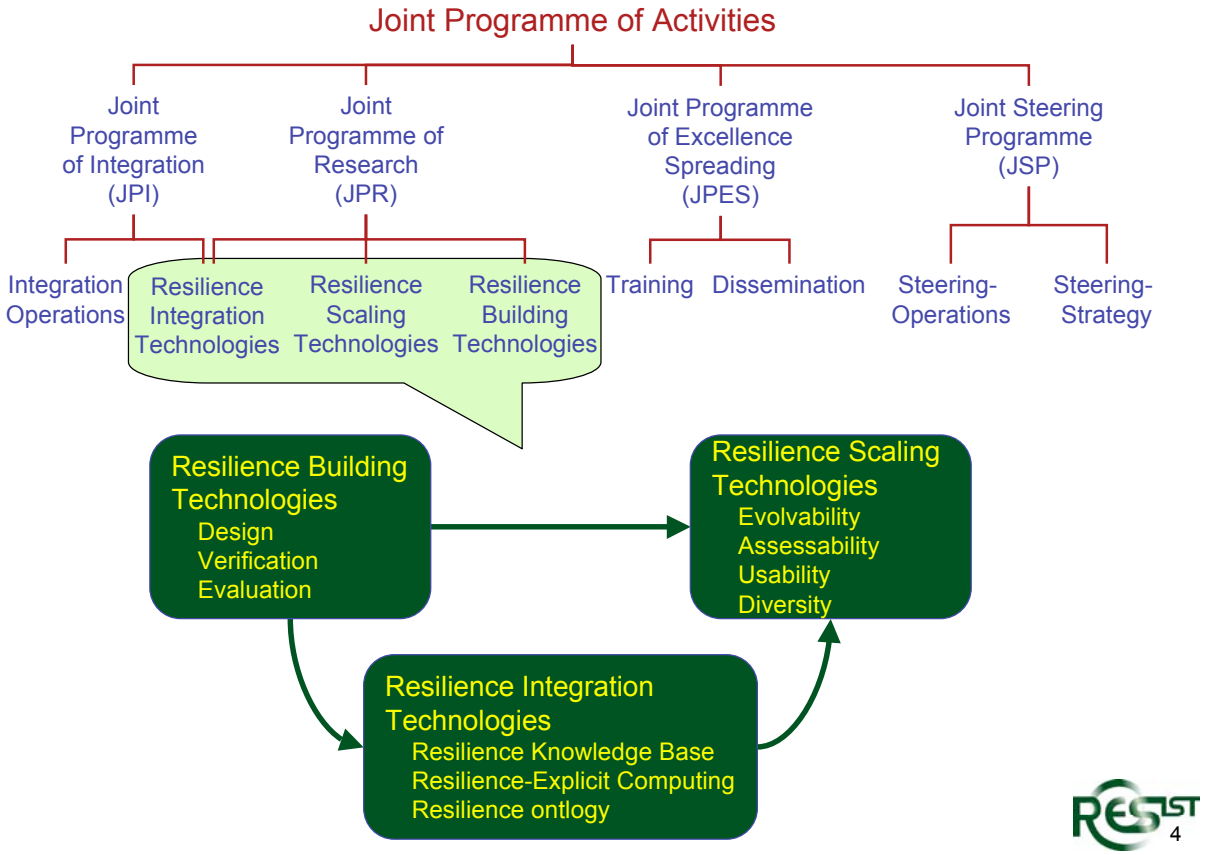

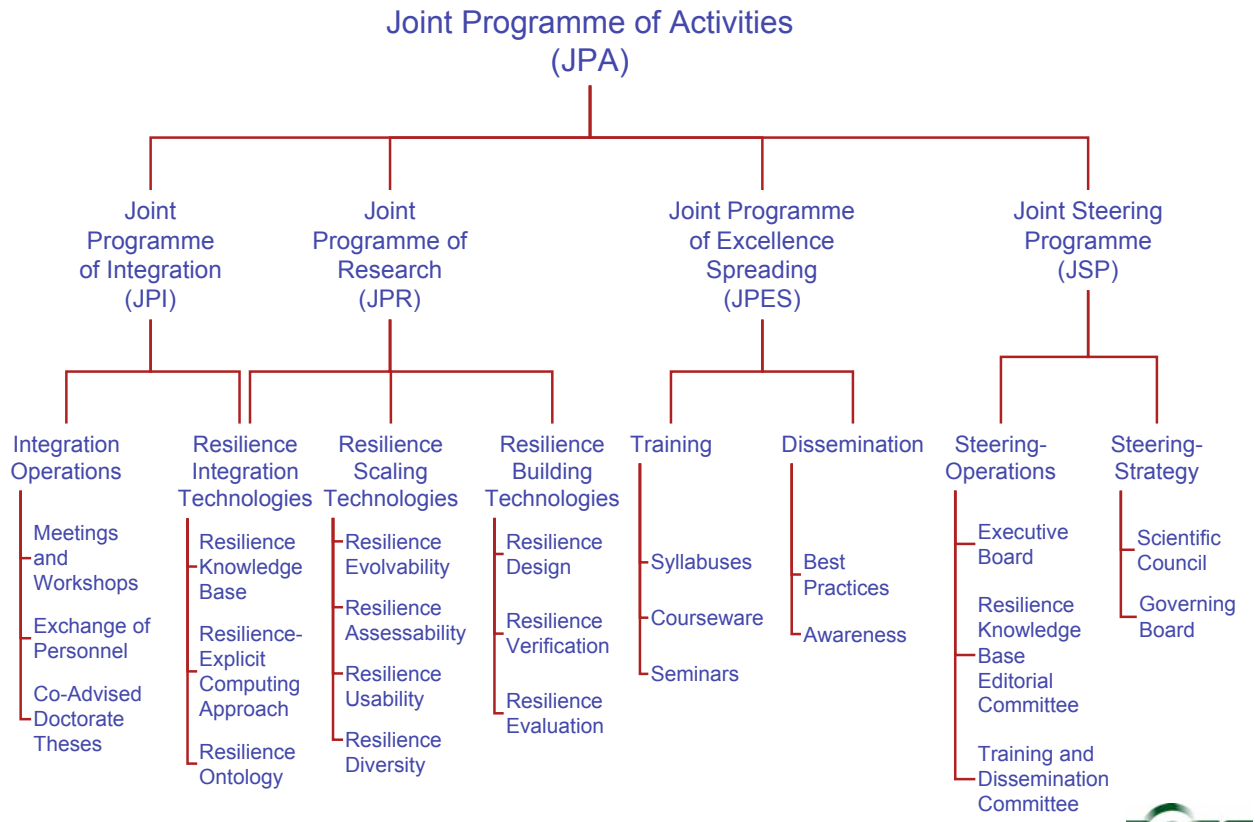

Partnership

110 researchers
61 students

	Expertise				Country	Academia (Ac) / Industry (Ind)
	Threat resilience: development or physical Accidental faults (A) / Malicious attacks (M) / Interaction mistakes (I)			Mobile computing		
	A	M	I			
LAAS-CNRS [coordinator]	X	X		X	FR	Ac
Budapest U.	X				HU	Ac
City U., London	X	X	X		UK	Ac
Darmstadt U.	X	X			DE	Ac
Deep Blue			X		IT	Ind - SME
Eurecom		X		X	FR	Ac
France Telecom R&D	X	X		X	FR	Ind
IBM Research Zurich		X			CH	Ind
IRISA	X			X	FR	Ac
IRIT			X		FR	Ac
Vytautas Magnus U., Kaunas	X				LT	Ac
Lisbon U.	X	X		X	PT	Ac
Newcastle U.	X	X	X		UK	Ac
Pisa U.	X	X	X		IT	Ac
QinetiQ	X	X			UK	Ind
Roma-La Sapienza U.	X			X	IT	Ac
Ulm U.	X				DE	Ac
Southampton U.	Resilience Knowedge Base building				UK	Ac

Organisation

☞ JPA - Workpackages

☞ Management

☞ Event Schedule

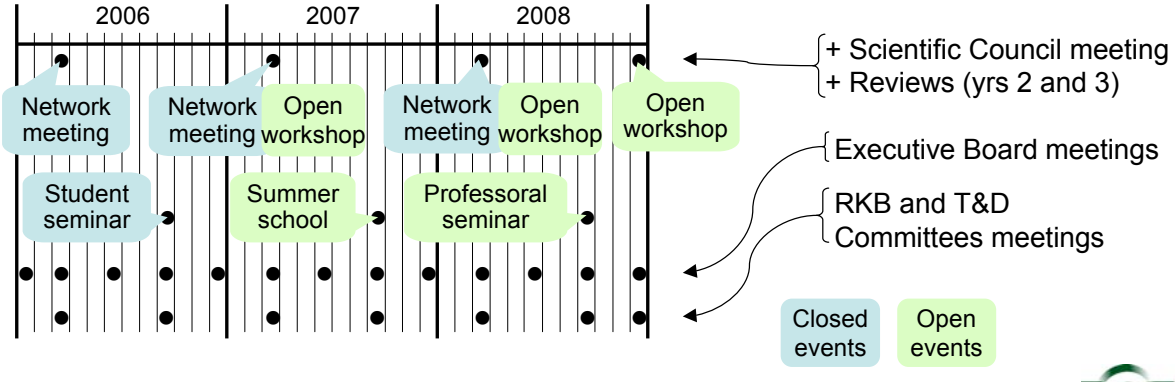

☞ Workpackages and organisational entities

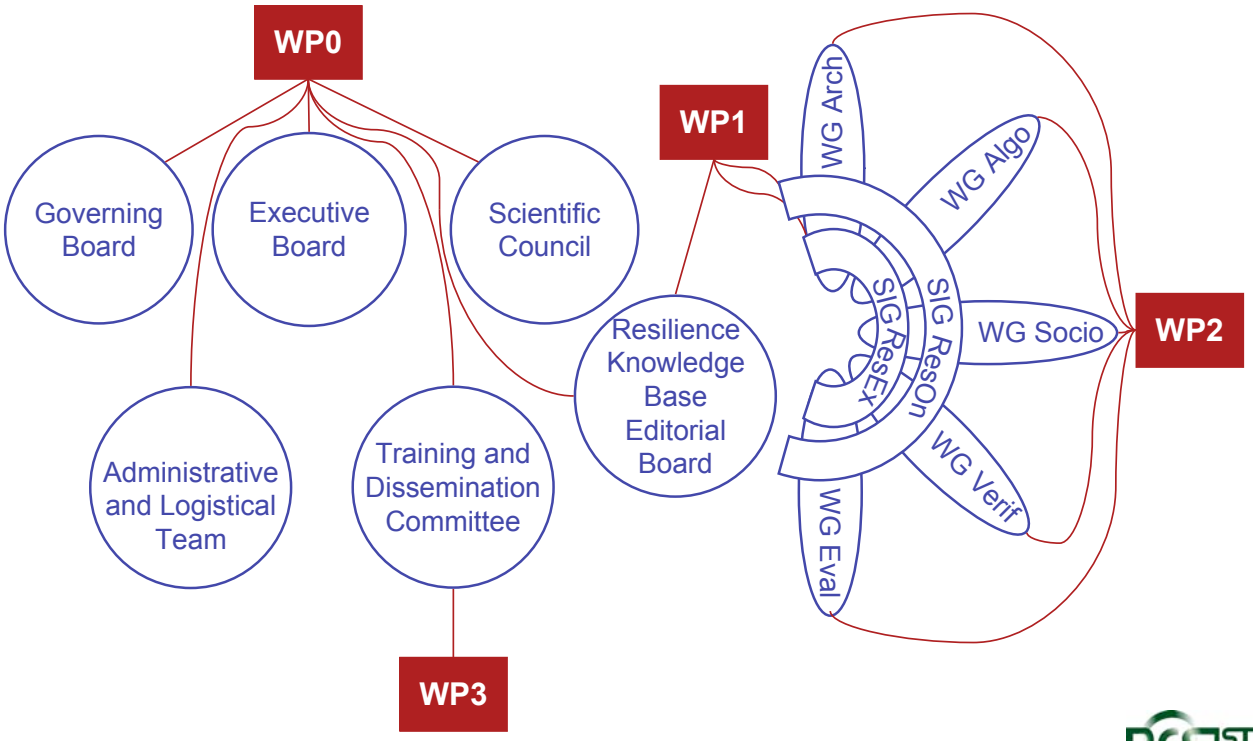

9

☞ Milestones

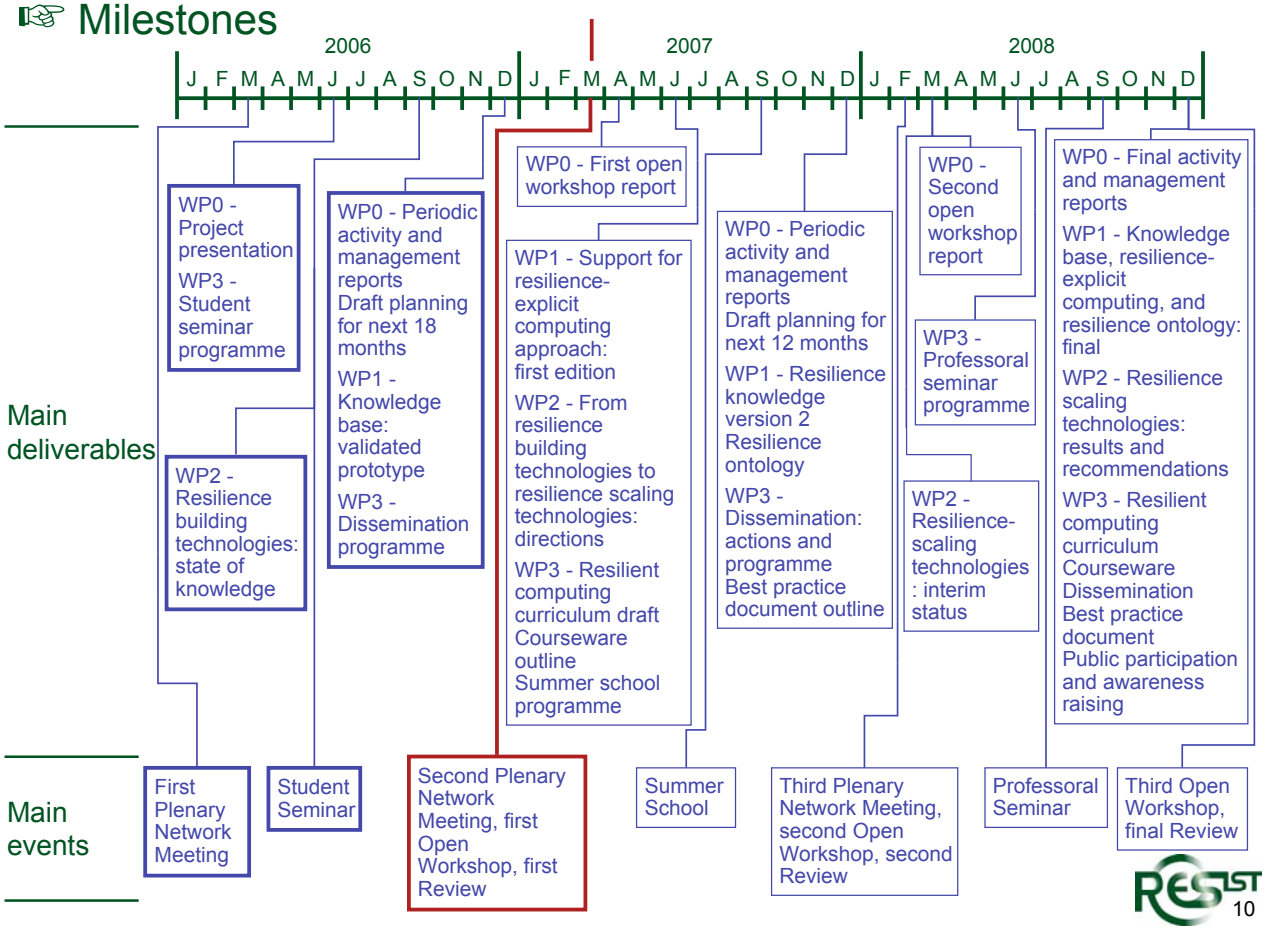

10

☞ Milestones

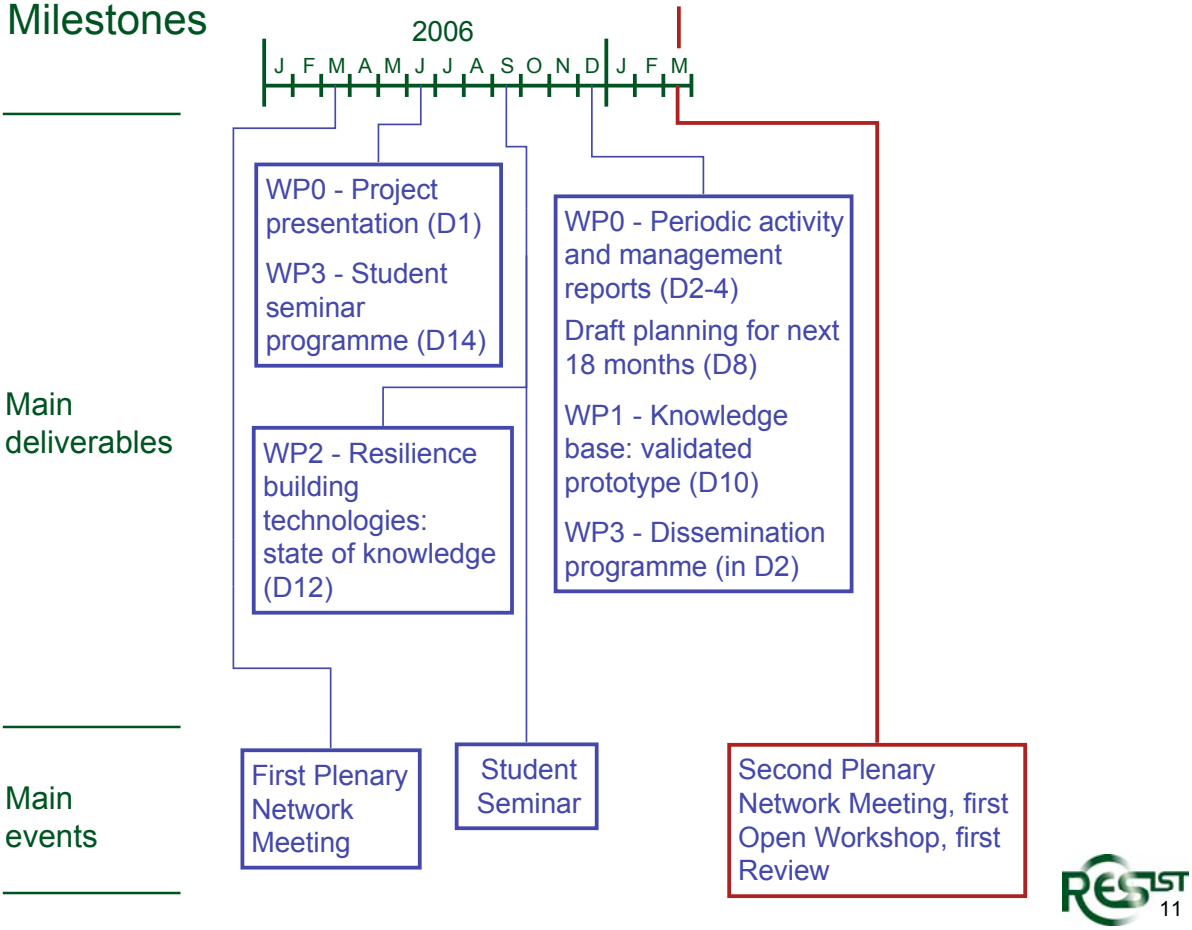

First year results

☞ Main Achievements

- ❖ State of Knowledge in Resilience-Building technologies
  - ➢ Main body
    - 5 parts (one per WG), 22 survey chapters
    - 68 co-authors from all ReSIST partners (54 researchers, 14 doctorate students)
    - Extensive review process, with emphasis on viewpoint of scientists who are not specialists of the sub-disciplines covered
    - A stepping stone in the process of integration
    - Substantial surveys that will be useful for the community at large
  - ➢ Appendices: Papers produced by ReSIST since January 2006

❖ Prototype Resilience Knowledge Base

- ➢ A semantic web environment for effective access to a body of knowledge on resilience concepts, methods and tools
- ➢ Current prototype: three classes of information, totaling 40 millions basic facts
  - ▪ Partners' resilience data
  - ▪ External sources including CORDIS, NSF, Citeseer, ACM publications, RISKS
  - ▪ Two ontologies: Dependability and Security, Systems concepts
- ➢ Information access enables relationships between entities to be displayed in the form of Communities of Practice
- ➢ Prototype reviewed by all ReSIST partners, and updated in response to feedback

☞ Significant events and advances

- ❖ Initial plenary meeting of the network (LAAS, 21-23 March), 101 ReSIST participants
- ❖ Student Seminar (San Miniato, Italy, 5-7 September), 32 Doctorate Students and 15 Senior Members
- ❖ Personnel exchange for at least one month stays, 5 ReSIST members, totalling 17 months of stay
- ❖ Co-advising of 4 doctorate theses.
- ❖ Production of 8 articles in scientific journals, and presentation of 52 communications (texts in proceedings)
- ❖ Presentation of ReSIST at 11 national, European and international events.

☞ Preparatory gound work

- ❖ Coming events, esp.
  - ➢ Open Workshop
  - ➢ Summer School, 24-28 September 2007, Porquerolles island
- ❖ Deliverables
  - ➢ Research Agenda, *From Resilience-Building to Resilience-Scaling Technologies: Directions*
  - ➢ Resilience-Explicit Computing Approach
  - ➢ Best Practice Document
  - ➢ Curriculum in Resilient Computing

Open Workshop and Review

- ❖ Salient results of the first year of activity
  - ➢ Selection of topics from the State of Knowledge document, covering all five WGs
  - ➢ Demonstration of the ontology-based resilience knowledge base
  - ☞ Comments, criticisms, and suggestions for future investigation welcome and expected
- ❖ Invited talks by two distinguished and highly renowned speakers
- ❖ Panel for presentation of resilience views by selected European projects (DESEREC, ESFORS, HIDENETS, SERENITY), and their comparison with ReSIST's views

About Resilience

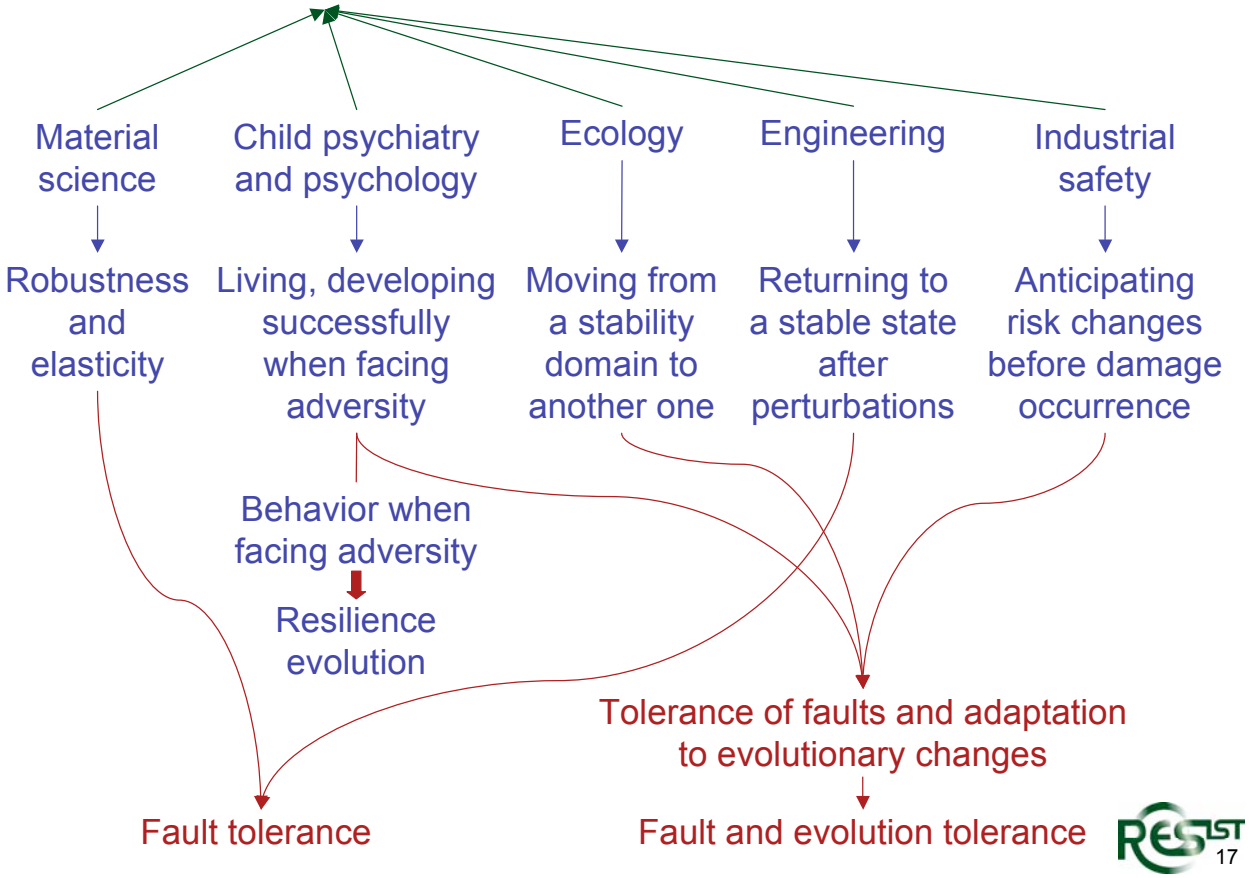

Computing systems and information infrastructures

- Resilience: ability to deliver, maintain, improve service when facing threats and evolutionary changes
  - Accidental and deliberate (esp. malicious)
  - functionnal, environmental, technological (hardware and software)
    - short term, e.g., dynamicity, mobility
    - medium term, e.g., new versions, reconfigurations
    - long term, e.g., reorganisations
- Failure: lack of adaptation to the (complexity of the) real world
  - Natural phenomena
  - Human-made features

1) Not (yet) a definition: evolutions ➠ threats
2) ⇒ « Re-visit » of the basic concepts of dependability
   - Extension of underlying system life-cycle model

MIDLAB Middleware Laboratory

Università di Roma "La Sapienza"
Dipartimento di Informatica e Sistemistica

Data Distribution in Large-Scale Distributed Systems

Roberto Baldoni
MIDLAB Laboratory
Università degli Studi di Roma "La Sapienza"

ReSIST: Resilience for Survivability in IST

First Open Workshop

Budapest 21-3-2007

SIXTH FRAMEWORK PROGRAMME

Information Society Technologies

What is a Large-Scale Distributed System?

What is a large-scale distributed systems?

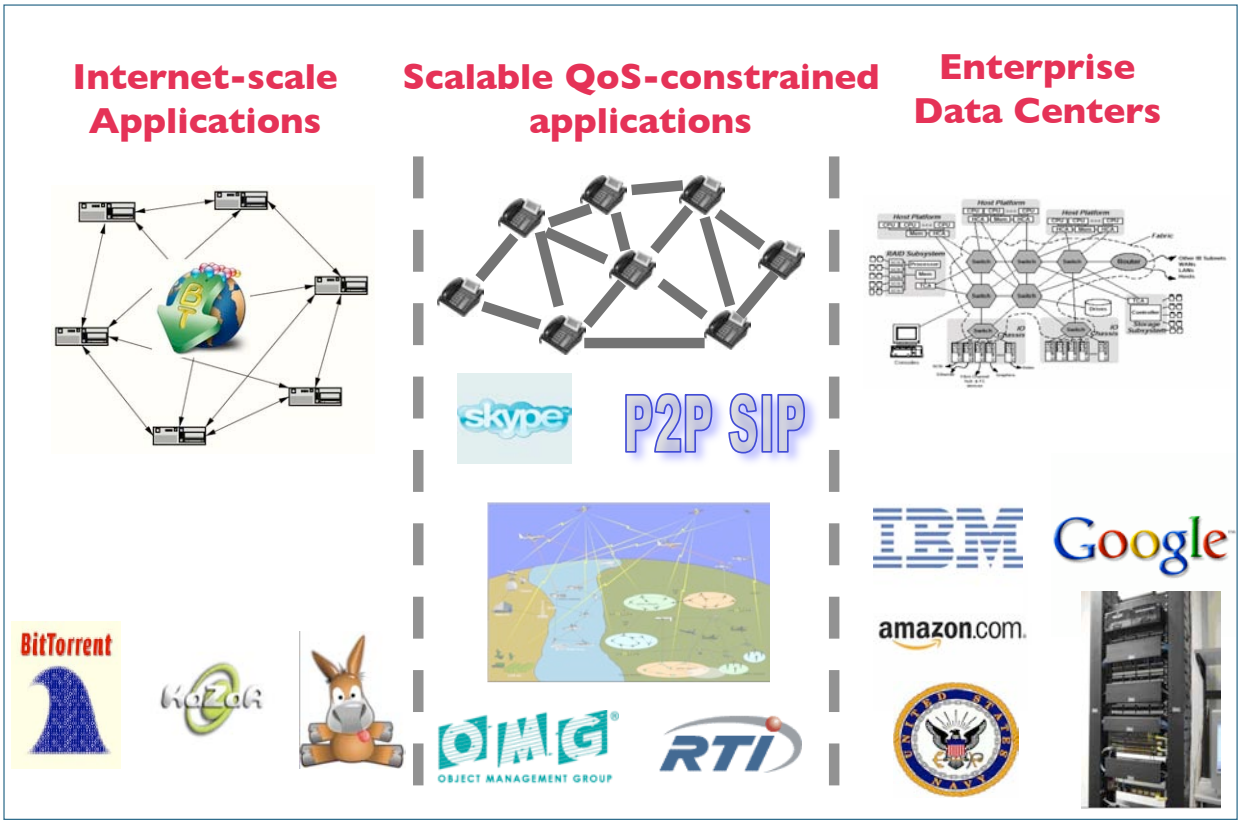

What is a large-scale distributed systems?

Internet-scale Applications
- unmanaged environment
- Shortlife peers
- High churn

Enterprise Data centers
- managed environment
- longlife peers
- low churn

Scalable QoS-Constrained Applicatio
- partially managed environment
- shortlife peers at network edges, longlife peers in the core
- high churn only at network edges, low churn in the core

Resilience while Scaling

MIDLAB Middleware Laboratory

What is the ideal software substrate for Large-Scale Distributed Systems?

P2P systems based on overlay networks

P2P systems based on overlay networks

Each application has requirements that impact the design of the overlay

Overlay Networks Substrate as superimposion of graphs

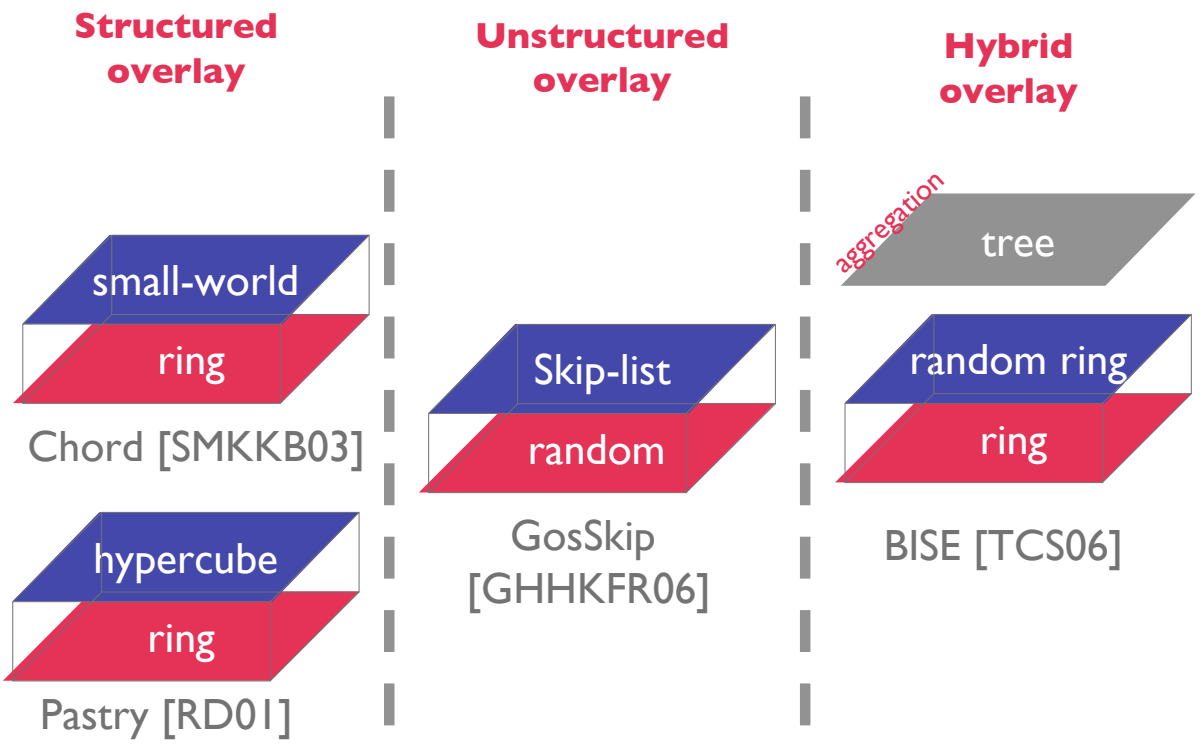

Middleware Laboratory
MIDLAB

Using publish/subscribe systems for Data Dissemination

Publishers: produce data in the form of **events**.

Subscribers: declare interests on published data with subscriptions.

Each **subscription** is a filter on the set of published events.

An **Event Notification Service** (ENS) notifies to each subscriber every published event that matches at least one of its subscriptions.

Interaction between publishers and a subscribers is **decoupled in space**, **time** and **flow**

- Two main models are considered in the literature

Topic-based selection

- Each event published in the system is tagged with a *topic* that completely characterizes its content.
- Each subscription contains a topic which the subscriber is interested in.

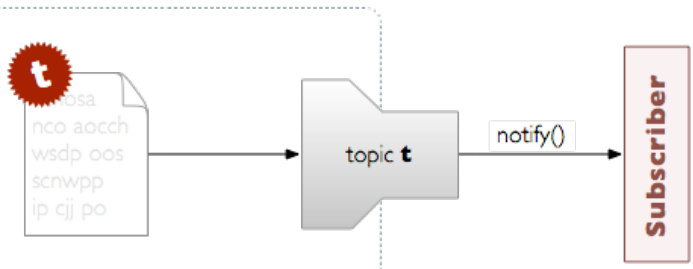

Content-based selection

- Each event published in the system is a collection of pairs *<attribute, value>*
- Each subscription is a conjunction of constraints over attributes.

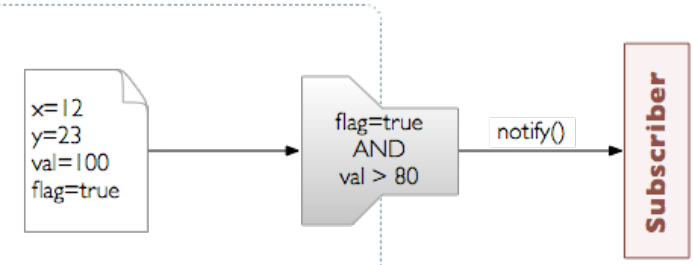

Scalable Data Distribution based on Overlay networks

Internet-scale Applications
- **Scribe [CDKR02], Pastry…**
- **Sub2Sub [VRKS06]**
- **TERA [BBQQVT07]**

Enterprise Data centers
- **BISE [TCS06]**
- **QuickSilver [OB07]**

Scalable QoS-constrained applications
- **Data Distribution Service (OMG)**
- **Control Plane (P2P SIP)**

Internet-Scale Data Distribution

- In a peer-to-peer environment peers play both the roles of publishers/subscribers and event brokers.
- Trivial solution to the problem of event dissemination:
  - Each event is broadcasted in the network.
  - Subscription-based filtering is performed locally.
- This usually implies a great waste of resources (on the network and on the nodes)
- The semantics of the publish/subscribe paradigm can be leveraged to confine the diffusion of each event only in the set of matched subscribers without affecting the whole network (traffic confinement)

MIDLAB Middleware Laboratory

Internet-Scale Data Distribution: Traffic confinement

- Traffic confinement can be realized solving three problems:
  - Interest clustering
    Subscribers sharing similar interests should be arranged in a same cluster; ideally, given an event, all and only the subscribers interested in that avent should be grouped in a single cluster.
  - Outer-cluster routing
    Events can be published anywhere in the system. We need a mechanism able to bring each event from node where it is published, to at least one interested subscriber.
  - Inner-cluster dissemination
    Once a subscriber receive an event it can simply broadcast it in the cluster it is part of.

Current solutions: Scribe

- Scribe [Castro et al., IEEE Journal on Selected Areas in Communications n.8 v.20, 2002]
  - Topic-based publish/subscribe implemented on top of DHTs.
  - For each topic a single node is responsible to act as a rendez-vous point between published events and issued subscriptions.
  - Problems:
    - single points of failure
    - hot spots
    - partial traffic confinement

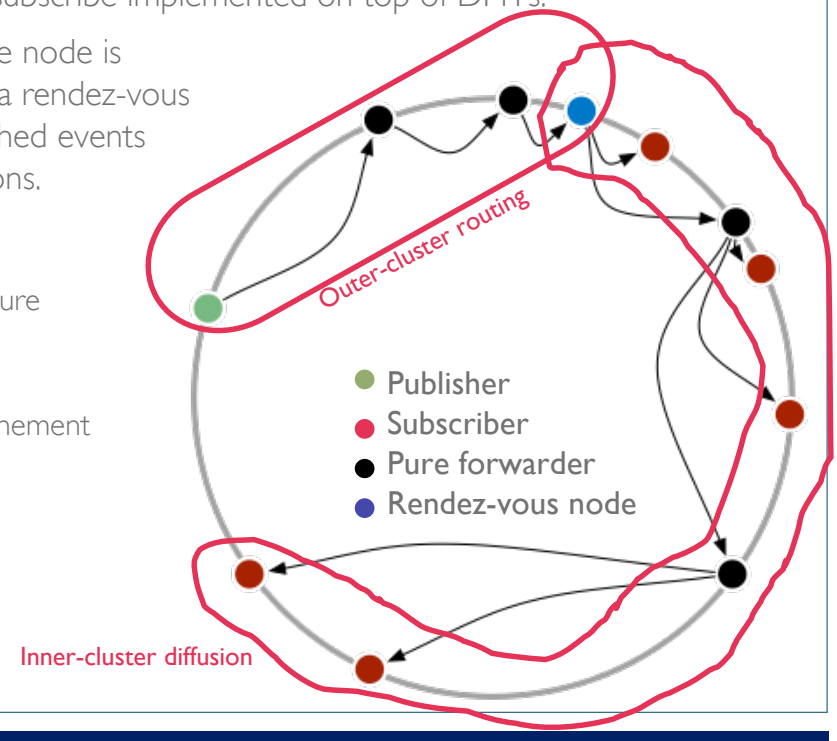

Current solutions: Sub-2-Sub

- Sub-2-Sub [Voulgaris et al., International Workshop on Peer-to-Peer Systems, 2006]
  - Content-based publish/subscribe
  - Complex three level infrastructure.
  - Employs clustering: brokers with similar interests are clustered in a same overlay.
  - Similarity is calculated checking intersections among subscriptions.
  - Problems:
    - depending on subscription distribution a huge number of distinct overlays must be maintained
    - the number of overlay networks a single node participates to is not proportional to the number of subscriptions it stores

Current solutions: Sub-2-Sub

- Sub-2-Sub [Voulgaris et al., International Workshop on Peer-to-Peer Systems, 2006]

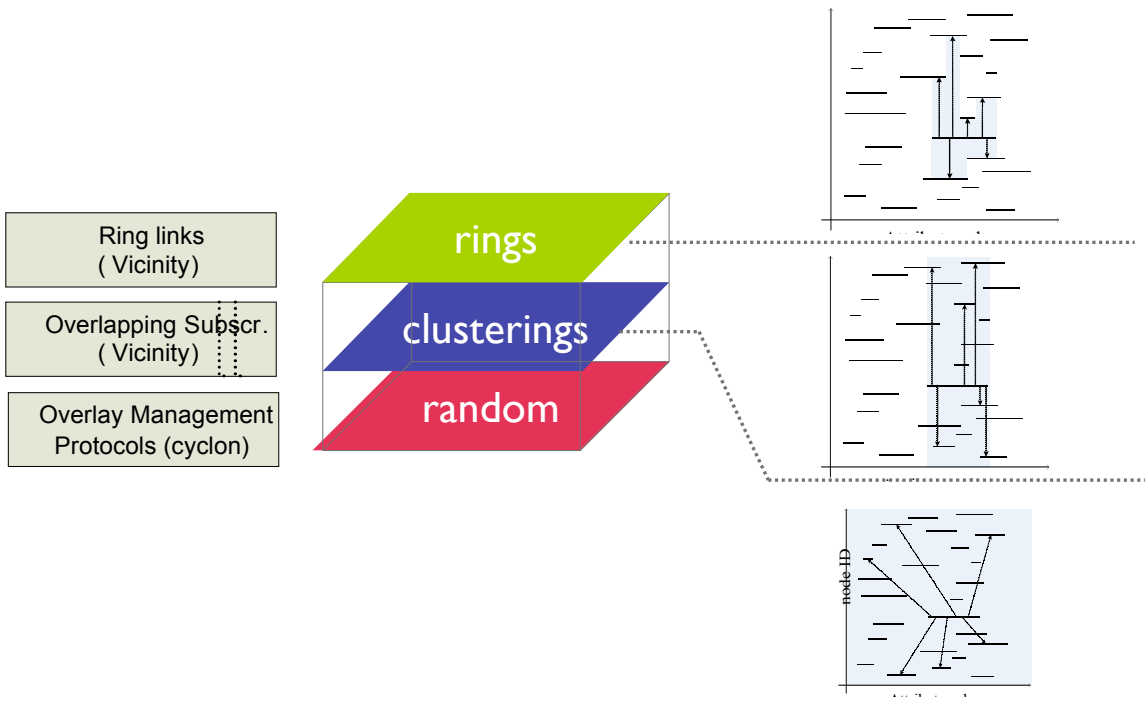

Middleware Laboratory
MIDLAB

TERA: Topic-based Event Routing for p2p Architecture

- A two-layer infrastructure:
  - All clients are connected by a single overlay network at the lower layer (general overlay).
  - Various overlay network instances at the upper layer connect clients subscribed to same topics (topic overlays).
- Event diffusion:
  - The event is routed in the general overlay toward one of the nodes subscribed to the target topic.
  - This node acts as an access point for the event that is then diffused in the correct topic overlay.

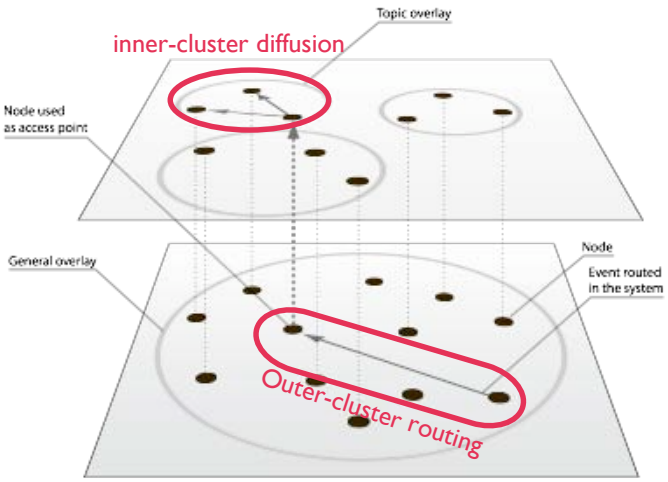

TERA: outer-cluster routing

- Event routing in the general overlay is realized through a random walk.
- The walk stops at the first broker that knows an access point for the target topic.

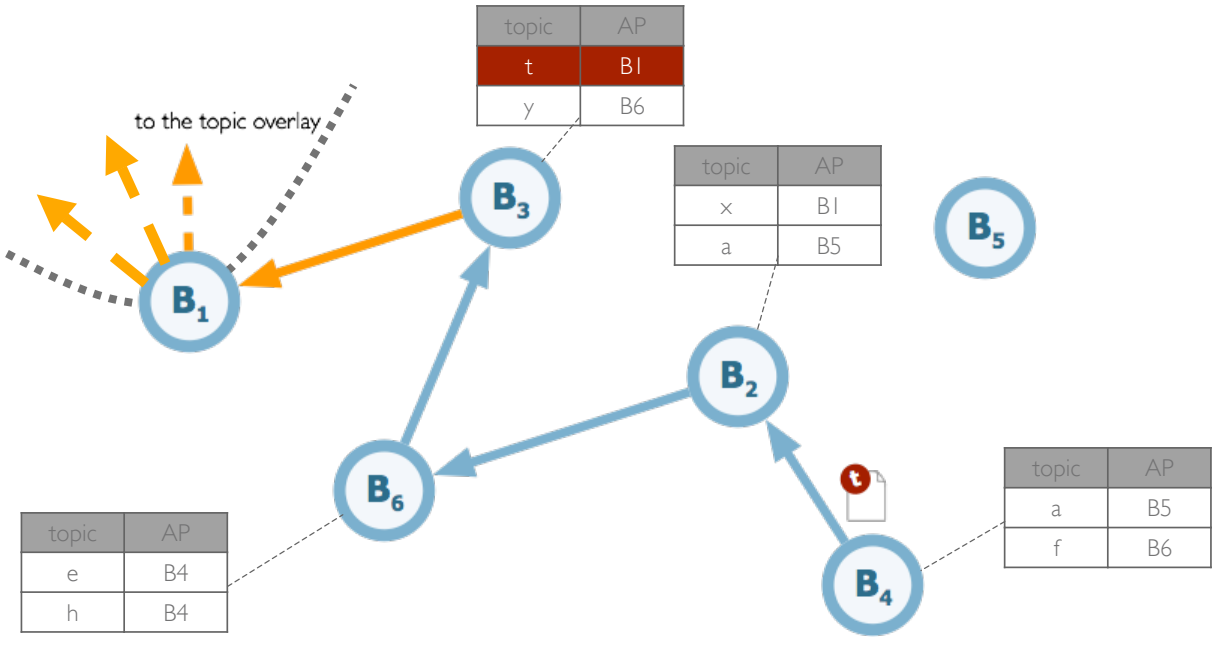

TERA: Architecture

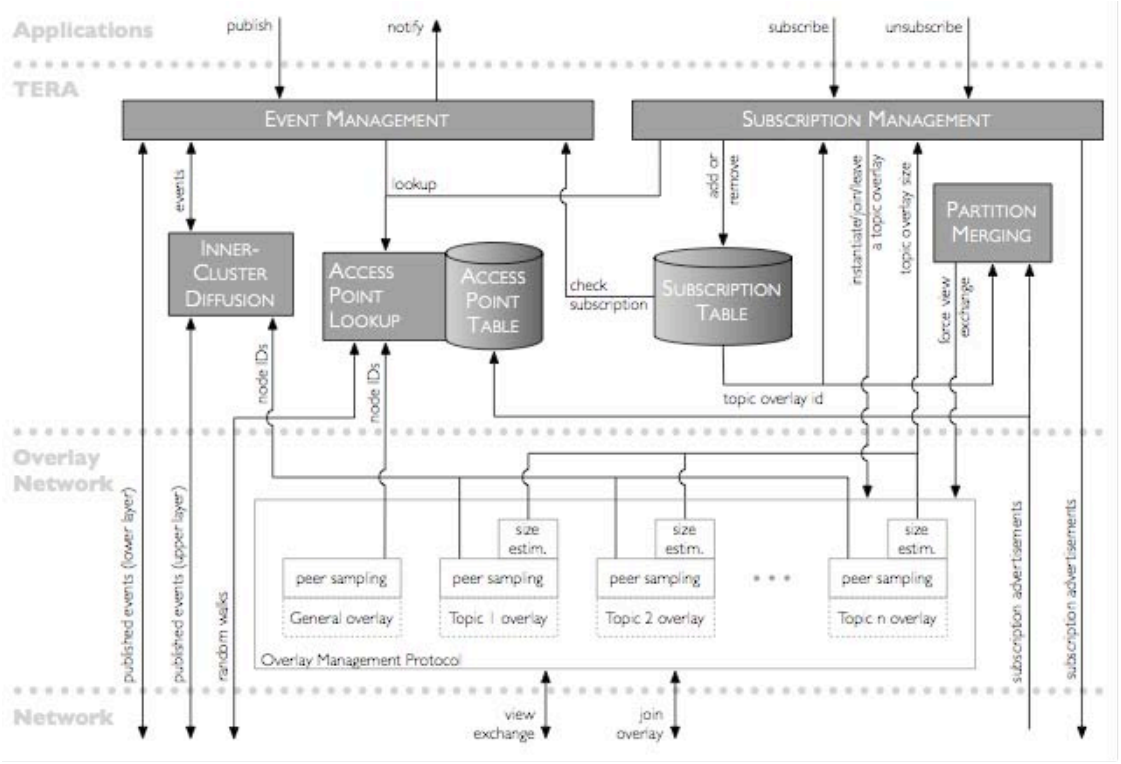

OMPs: Newscast, Cyclon, etc.

TERA Results: Outer-cluster routing

- We want every topic to appear with the same probability in every APT, regardless of its popularity.

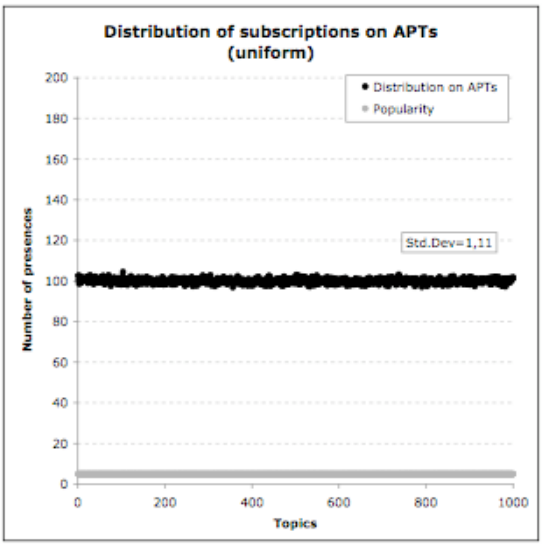

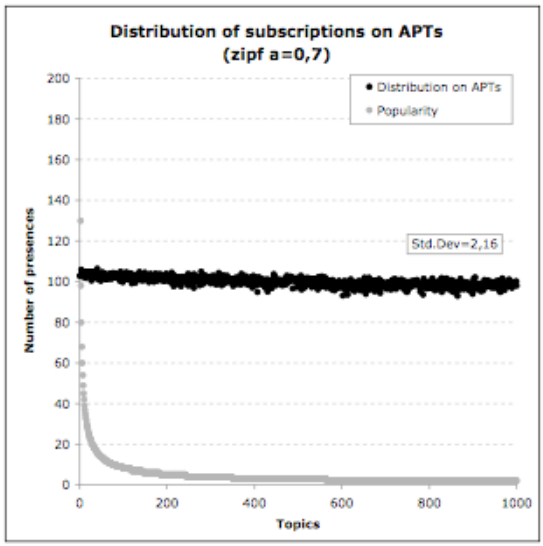

TERA Results: traffic confinement

Which is the probability for an event to be correctly routed in the general overlay toward an access point ?

- Depends on:
  - Uniform randomness of topics contained in access point tables.
  - Access point table size.
  - Random walk lifetime.

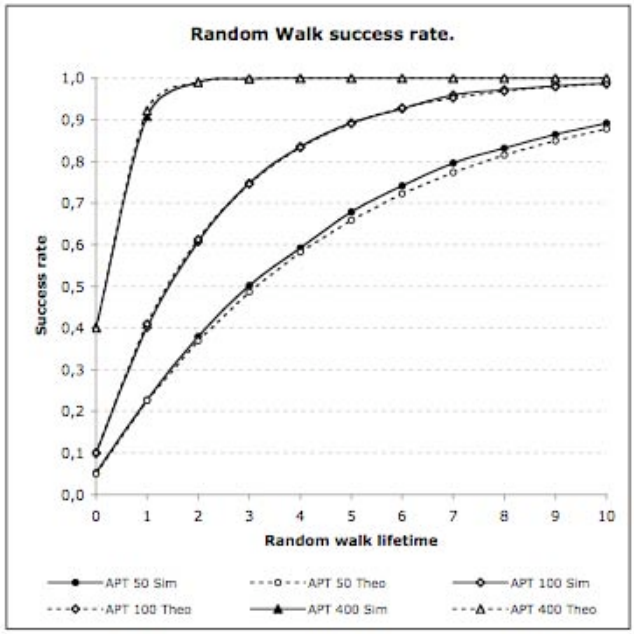

Conclusions

- Scalable Data Distribution based on Overlay networks for Internet-Scale applications
  - What is a large scale distributed systems
  - P2P Overlay networks as the ideal substrate for
    - Internet-scale applications
    - Enterprise datacenter applications
    - Scalable QoS-constrained applications
- TERA: Topic-based Event Routing for p2p Architecture
  - outer-cluster routing
- Joint activities within RESIST
  - Composing gossiping: a conceptual architecture for designing gossip-based applications. R. Baldoni, H. L, J. Pereira, É. Rivière (Submitted paper)
  - A Component-based Methodology to Design, Arbitrary Failure Detectors for Distributed Protocols. R. Baldoni, J.M. Helary, S. Tucci Piergiovanni. ISORC 2007
  - Looking for a Definition of Dynamic Distributed Systems. R. Baldoni, M. Bertier, M. Raynal, and S. Tucci-Piergiovanni (submitted paper)

MIDLAB Middleware Laboratory

Cooperative Backup for Dynamic Systems

- Dynamic Systems in a Ubiquitous World
  - Nomadic devices
  - Mostly disconnected operations
  - Opportunistic wireless communication with similar devices
  - Peer-to-peer model of interactions
  - Embedded data generation
- Secure Cooperative Backup for Nomadic Devices
  - Leverage encounters for storing data
  - Even when no infrastructure is available

Cooperative Backup for Dynamic Systems

- Backup = protection of critical private data against
  - Permanent and transient faults affecting a data owner
  - Theft or loss of a data owner

Cooperative Backup for Dynamic Systems

- Backup = protection of critical private data against
  - Permanent and transient faults affecting a data owner
  - Theft or loss of a data owner
- New threats on backups
  - Malicious (and accidental) faults
  - Confidentiality, integrity and availability
- New threats on service
  - Selfish denial of service (refusal to cooperate)
    - Free-riding : consumption without contribution
    - “Tragedy of the commons” (Hardin 1968)
    - Attacks must be made unprofitable
  - Malicious denial of service (sabotage)
    - Attacks must be made ineffective or too costly

Cooperative Backup for Dynamic Systems

- Challenges
  - No prior organization
  - Ephemeral interactions
  - Limited energy, computation and storage
  - Only intermittent access to a fixed infrastructure
- Usual criteria for classic functionalities
  - User transparency
  - Usability
  - etc.

Overview

- Motivations
- Data Availability: Data scattering
  Data encoding and redundancy control [Courtès et al. 07]
  - (n,k) codes
  - Evaluation using GSPN and Markov chains
- Service Availability: Cooperation Incentives
  Crypto-challenges that can be delegated [Oualha et al. 07]
  - Probabilistic cooperation checking
  - Evaluation using game theory

Scattering and Redundancy

- Opportunistic communication to peers and to infrastructure
- Ephemeral encounters
  - Duration/frequency ?
  - Amount of data ?
  - Reliability of contributors ?
  - ➔ Scattering of fragments
- Untrusted and unreliable contributors
  - Ability to get fragments back ?
  - ➔ Replicate fragments
- Limited storage resources
  - ➔ Trade-off between redundancy and resource use
  - ➔ Optimization of gained availability vs resources

➔ Modeling and evaluation of scattering policies

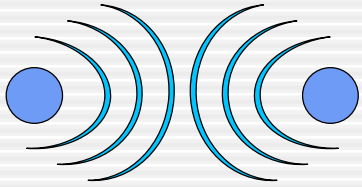

Examples

Classic redundancy

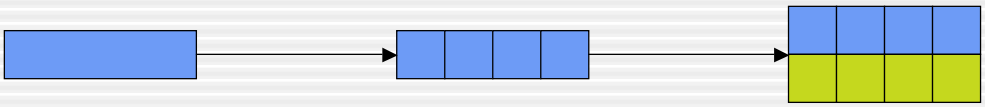

➔ 1 fault / size = x2

(n,k) codes

➔ 4 faults / size = x2

(n,k) codes

(n,k) codes

Input: k fragments

Output: n fragments

➔ n-k faults / size = n/k

Which code to choose?
How to choose N and K?

→ Modeling and Evaluation

GSPN model for (n,k) erasure codes

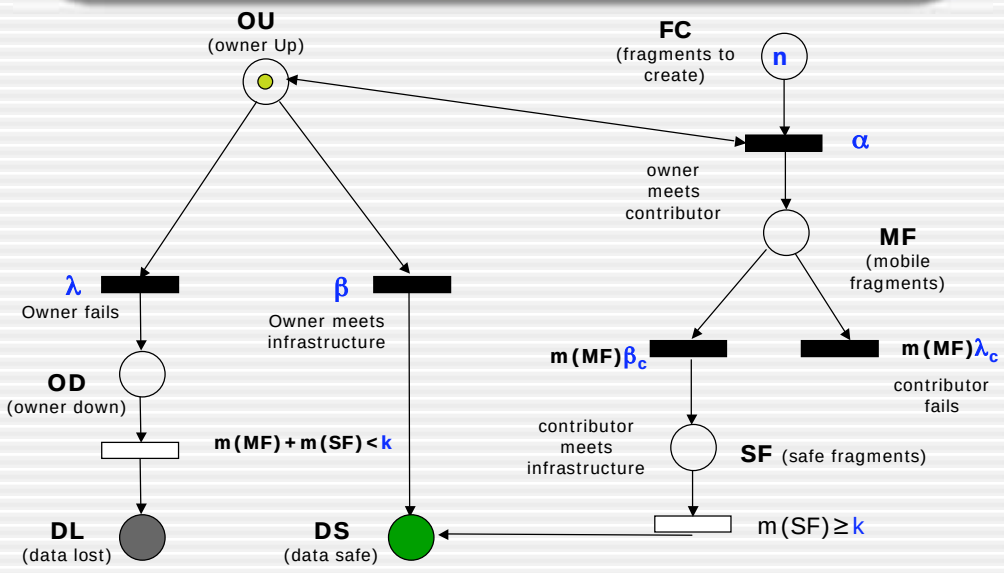

Sensitivity analyses: summary

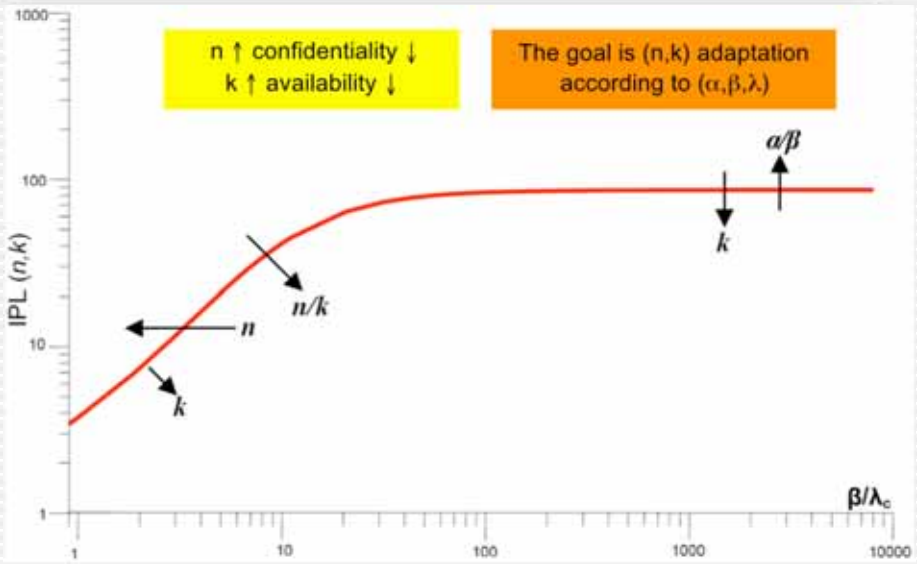

Service Availability

- Resource sharing
  - "Tragedy of the Commons" [Hardin68]
  - Free-riding (consumption without contribution)
- *Cooperation incentives*
  - *Money* (e.g., Buttyan's nuglets, claims, etc.)
    - Trade money for service
  - *Reputation*
    - Detect misbehavers, give them bad reputation
    - Don't cooperate with devices with bad reputation

Buttyan's nuglets

- Each node maintains a counter (nuglet)
  - Decreased when sending its own packet
  - Increased when forwarding a packet
  - The counter must remain positive

- The policy must be enforced
  - Use of tamperproof hardware
    - SIMcards, JavaCards, etc.
    - TPM

Marti's Watchdogs

- Each node possesses a watchdog
  - When a node sends a packet, the watchdog verifies that the neighbors forward it

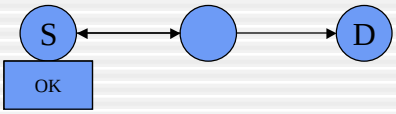

Marti's Watchdogs

- Each node possesses a watchdog
  - When a node sends a packet, the watchdog verifies that the neighbors forward it

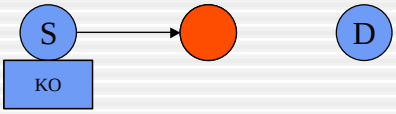

- Misbehaving nodes are detected:
  - Bad reputation
  - No cooperation

Reputation Establishment

- Reputation has to be based on cooperation observation
  - Does a contributor contributes ?
- Cooperative backup: does a contributor stores the data ?
  - Test it with challenges
  - Long-term and disconnected service
  - Challenges have to be delegated

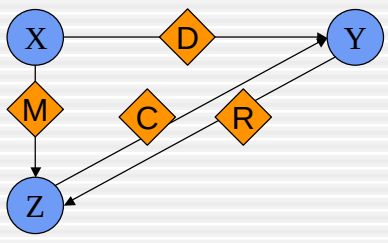

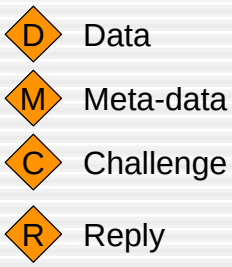

Reputation Establishment

- Crypto-challenges that can be delegated [Oualha et al. 07]
  - Probabilistic verification
    - D — Data = Signed data
    - M — Meta-data = Public Key + # blocks
    - C — Challenge = Random block id
    - R — Reply = Signature of chosen block
- Z verifies the challenge reply to
  - Establish Y reputation
  - Choose to cooperate with Y

Current and Future work

- More general evaluation assumptions
  - Trust and cooperation wrt participating nodes (malicious, selfish)
  - Other dissemination strategies
- Adaptable Scattering Strategy
  - Online evaluation of (α,β,λ)
  - According to the user preferred policy
  - Compute and apply the best strategy
- Cooperative geo-service providing
  - A service is associated to a path
  - Nodes in the vicinity of the path cooperate to provide the service
- Failure detectors targeting cooperation faults
  - DoS attacks, Sybille attacks, etc.

Information Society Technologies

SIXTH FRAMEWORK PROGRAMME

Challenges and Advances in E-voting Systems

Technical and Socio-technical Aspects

Peter Y A Ryan

Lorenzo Strigini

Outline

- The problem.
- Voter-verifiability.
- Overview of “Prêt à Voter”.
- Resilience and socio-technical aspects
- Conclusions.
- Future work (in ReSIST)

The Problem

- Highly adversarial: system trying to cheat voters, voters trying to cheat the system, coercers trying to influence voters, voters trying to fool coercers etc.
- The Ancient Greeks experimented with primitive technological solutions to try to shift the trust from people (officials) to mechanical devices.
- In the US technological devices for voting have been used for over a century: e.g., lever machines since 1887, punch cards, optical scans, touch screen etc. prompted by high instance of fraud with paper ballots!
- All have problems, see “Steal this Vote” Andrew Gumbel.

“The Computer Ate my Vote”

- In the 2004 US presidential election, ~30% of the electorate used DRE, touch screen devices.
- Aside from the “thank you for your vote for Kerry, have a nice day” what assurance do they have that their vote will be accurately counted?
- What do you do if the vote recording and counting process is called into question?
- Need to trust the (proprietary) software.
- Voter Verifiable Paper Audit Trail (VVPAT) and “Mercuri method” have been proposed. But paper trails are not infallible either.
- Nedap machines in the Netherlands etc.

Florida 2000

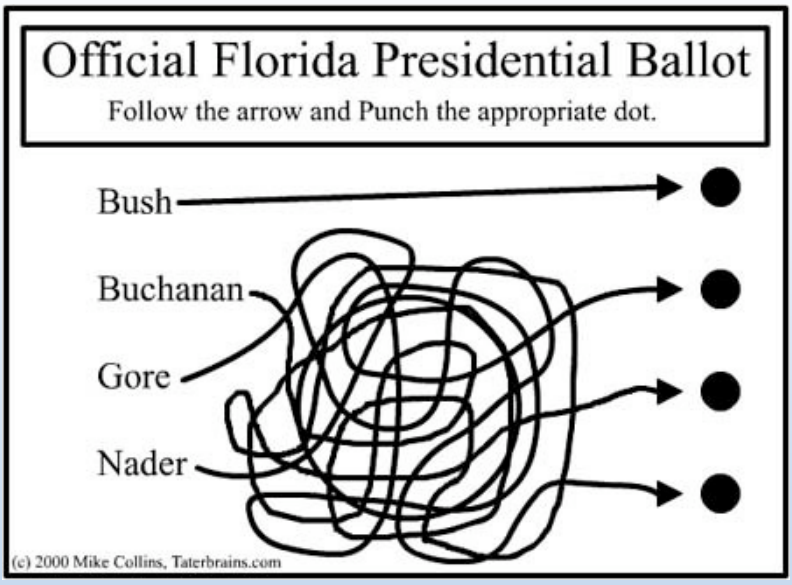

The challenge

- Digital voting technologies hold out promise of accessible and efficient democracy.
- Want high assurance that all votes are accurately recorded and counted-whilst maintaining ballot secrecy.
- The challenge is to reconcile these two conflicting requirements whilst minimising, ideally eliminating, dependence on the components (devices, tellers, software, hardware, officials etc.) of the scheme.
- Needs to be usable and sufficiently understandable to be widely trusted.

Technical Requirements

- Elections should be “free and fair”.
- Typical, key requirements:
  - (unconditional) integrity: count accurately reflects votes cast.
  - Ballot secrecy: the way a voter cast their vote should only be known to the voter.
  - Voter verifiability: the voter should be able to confirm that their vote is accurately included in the count and prove to a 3rd party if it is not (without having to revealing their vote).
  - Universal verifiability: anyone should be able to verify the count.
  - Availability: all eligible voters should be able to cast their vote without let or hindrance throughout the voting period.
  - Ease of use, public understanding and trust, cost effective, scalable etc. etc…..

Assumptions

- For the purposes of the talk we will make many sweeping assumptions, e.g.:
  - An accurate electoral register is maintained and available.
  - Mechanisms are in place to ensure that voters can be properly authenticated.
  - Existence of a secure Web Bulletin Board.
  - Crypto algorithms are sufficiently secure.
  - Etc.

Voter-verifiability in a nutshell

- Voters can confirm that their vote is accurately but not prove to a third party how they voted.
- Voters are provided with an encrypted “receipt”.
- Copies of the receipts are posted to a secure web bulletin board. Voters can verify that their (encrypted) receipt is correctly posted.
- A (universally) verifiable, anonymising tabulation is performed on the posted receipts.
- Checks (random audits) are performed at each stage to detect any attempt to corrupt the encryption and the decryption or the receipts.
- The guarantees of integrity are not dependent on correct behaviour of software, hardware, officials etc.

Voting with commuting diagrams

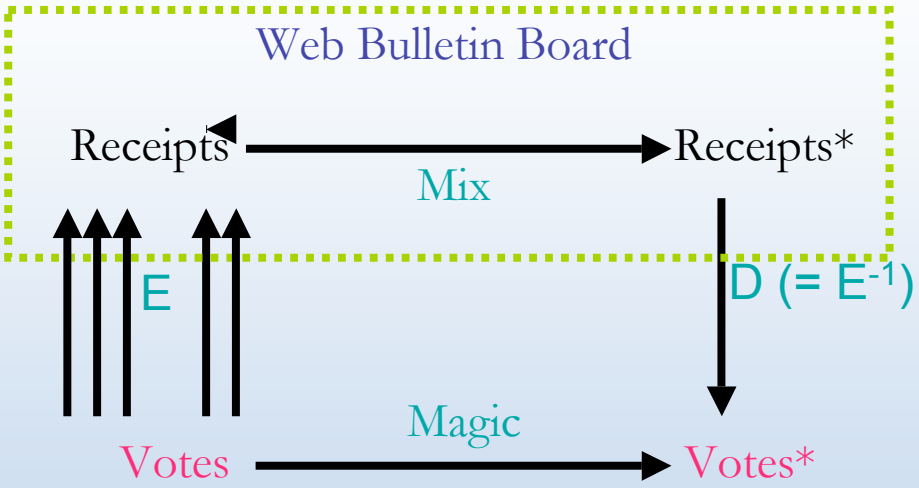

Prêt à Voter

- The key innovation of Prêt à Voter is to encode the vote by randomising the candidate order.
  - Voter experience simple and familiar.
  - Votes are not directly encrypted, just the frame of reference in which votes encoded. Hence:
    - The vote recording device doesn't get to learn the vote.
    - No need for ZK proofs of correct encryption of votes-but onus of proof shifts to showing the well-formedness of the ballot forms.
    - Avoids subliminal, kleptographic and side channels.
- Prior work: Chaum, Benaloh, Neff,…

Typical Ballot Sheet

Obelix	
Asterix	
Idefix	
Panoramix	
Geriatrix	
	$rJ9*mn4R&8

Voter marks their choice

Obelix	
Asterix	x
Idefix	
Panoramix	
Geriatrix	
	$rJ9*mn4R&8

Voter's Ballot Receipt

	Cast-valid
x	
$rJ9*mn4R&8 449034729948	

After the voting phase

- Once the election is closed, digital copies of the receipts are posted to the Web Bulletin Board (WBB).
- The voters can visit the WBB and confirm that their receipt appears correctly.
- Additionally, checks could be performed by independent entities between the (encrypted) paper audit trail and posted receipts.
- A verifiable, anonymising tabulation is performed with all intermediate stages posted to the WBB.

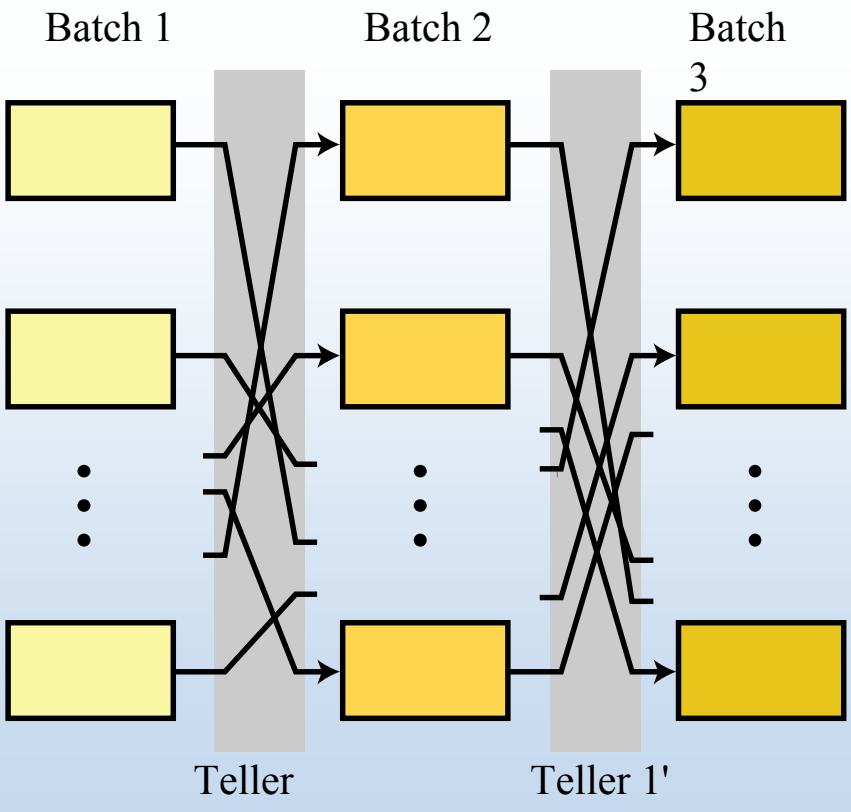

Auditing the tellers

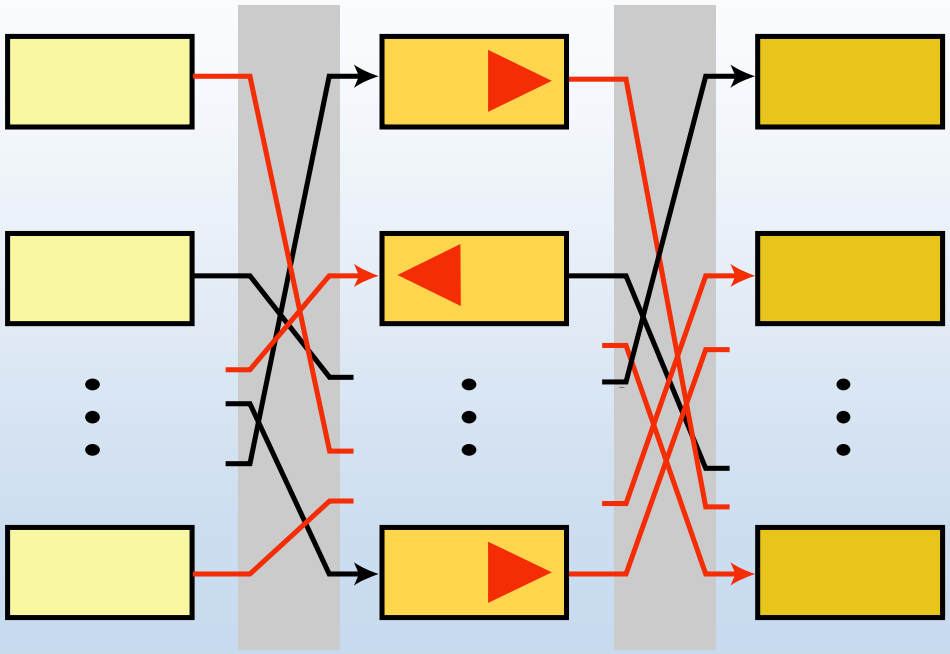

Enhancements

- Vulnerability analysis.
- Randomising encryption and re-encryption mixes.
- Distributed generation of encrypted ballots.
- On-demand decryption and printing of ballot forms.
- (A variant of) Adida/Rivest off-line audit mechanism.
- Coercion-resistant remote variants (with Cornell).
- Crypto-free, scratch card version.

Resilience aspects

- cryptography-supported voter-verifiability promises much
  - *more* integrity and privacy than paper systems
  - run-time monitoring reduces need for special, heavily verified machinery
- but there is more to a voting *system*
  - error/attack detection does not make error/attack tolerance
  - .. recovery delegated to human part of system

ICT fault tolerance in the election system

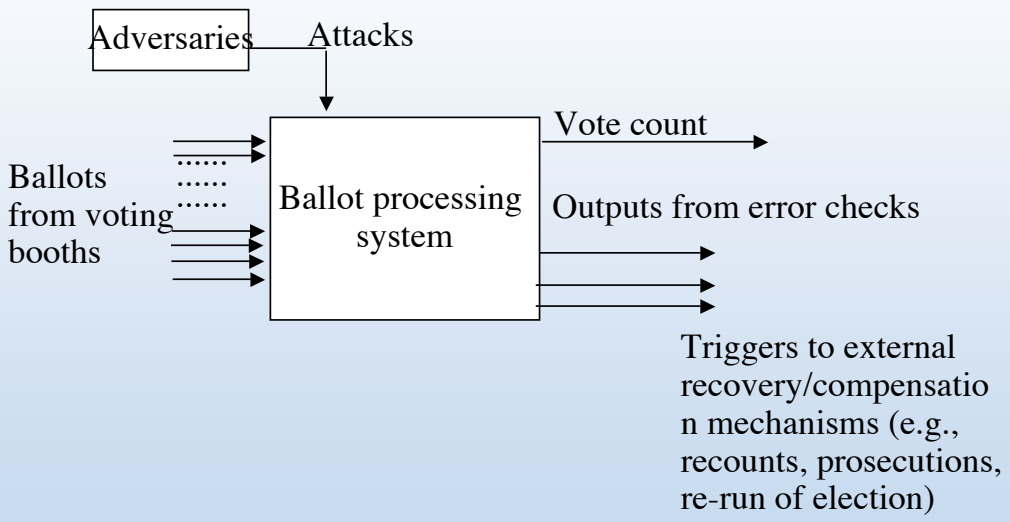

Effects of strong error detection

- election corruption is made more difficult
- but detected errors are expensive, so:
  - error recovery (automated and human) is important
  - better coverage may shift attackers' preference, e.g. from attempting *undetected* vote corruption to simply sinking the election
  - good integrity and privacy; *availability* issues
    - e.g. DDoS attacks on bulleting boards?
    - increased requirements for ICT support to be robust/resilient

Wider socio-technical aspects

- attacker's target might become simply the *reputation* of the election system
- implications cross the boundary between what can be designed (hardware, procedures) and political management
- so, a range of issues
  - from user-friendliness, HCI of voting machines
  - to choice of algorithms that public will be able to trust
  - to ensuring enough parties do perform the checks that anyone *may* perform
  - to ensuring *correct* perception of trustworthiness of each specific election

Conclusions

- we have presented: a technical problem, some solutions
  - Maximal transparency (consistent with ballot secrecy).
  - Accuracy independent of software, hardware, etc.
  - High assurance of detection of corruption.
  - Verify the election not the system!
- And open issues

Conclusions cont.

- E-voting is a ReSIST problem par excellence..
  - large distributed system, complex dependability requirements, evolving threats
  - “must work well the first time around”, *every* time - implying need for resilience
  - ICT entwined with users and their reactions

Future work

- Further enhancements (simplifications!?)
- Further analysis of the resilience of the system
- Investigate recovery mechanisms and strategies
- Investigate socio-technical aspects
- Investigate public understanding and trust
- Basis for a ReSIST case study

March 21, 2007 - ReSIST Open Workshop, Budapest, Hungary

Information Trust INSTITUTE

Probabilistic Validation of Computer System Security

William H. Sanders
University of Illinois
(Joint work with DPASA Project Team)
www.iti.uiuc.edu

Everyone says it is important, few approaches exist …

- Security metrics were an important problem in the 2005 INFOSEC Research Council Hard Problems List
- New security metrics that are linked to the business were ranked first among six key security imperatives developed by over twenty Fortune 500 firms
- New regulatory requirements of Sarbanes-Oxley and the Basel II Accord have created more urgency for metrics that integrate security risk with overall business risk
- Almost every critical infrastructure roadmap lists security metrics as a critical challenge
- The list goes on …

ITI

Security Validation Truths …

- Security is no longer absolute
- Trustworthy computer systems/networks must operated through attacks, providing proper service in spite of possible partially successful attacks
- Intrusion tolerance claims to provide this ability
- If security is not absolute, quantification of the "amount" of security that a particular approach provides is essential
- Quantification can be useful in:
  - A *relative* sense, to choose amount alternate design alternatives
  - In an *absolute* sense, to provide guarantees to users

Existing Security Validation Approaches

- Most traditional approaches to security validation have focused on and specifying procedures that should be followed during the design of a system (e.g., the Security Evaluation Criteria [DOD85, ISO99]).
- When quantitative methods have been used, they have typically either been based on:
  - formal methods (e.g., [Lan81]), aiming to prove that certain security properties hold given a specified set of assumptions, or
  - been quite informal, using a team of experts (often called a "red team," e.g. [Low01]) to try to compromise a system.

Problems with Existing Approaches

- Process Guidelines can improve security, but provide NO quantification of the amount of security that has been obtained
- Formal methods aim either to prove absolute security (not usually possible), or find problems (useful, but NO quantification.
- Red Teams, can find problems (useful), but again, no quantification (sample size too small).
- Most existing metrics are lagging indicators of performance (and hence not predictive!)
- Probabilistic Methods can provide predictive quantification, but their application to security/ survivability is challenging as well.

Security Quantification Challenges

- How can the behavior of attackers be quantified?
  - How accurately does this need to be done?
  - At what level of detail?
- How should security/survivability measures be specified?
  - Are new measures needed?
- If relative measures are desired, can they be shown to be robust across a wide variety of situations?
  - Robustness is key to good design
- How accurately can absolute measures be estimated?
- Can quantification aid in security testing?
  - Knowing where to focus testing is key
- Can a notion of "coverage" be developed?
  - If so, testing can produce quantitative results

Example Probabilistic Security Validation Study

- Evaluation of DPASA-DV Project design
  - Designing Protection and Adaptation into a Survivability Architecture: Demonstration and Validation
  - USA DARPA Project, 2.5 years; 11 Million $, ~25 people on project team.
- Design of a "Joint Battlespace Infosphere"
  - Publish, Subscribe and Query features (PSQ)
  - Ability to fulfill its mission in the presence of attacks, failures, or accidents
- Goal was to design AND validate survivability of system while operating under intense attack

ITI

JBI Design Overview

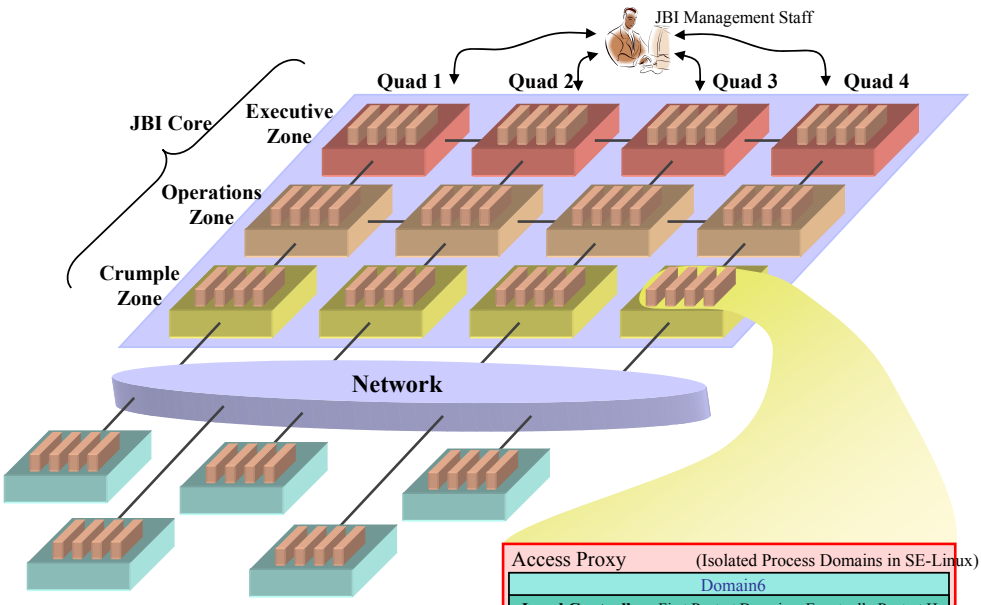

Protection Domains
Isolation among selected functions on individual core hosts and on clients

ITI

Survivability/Security Validation Goal

- Phase 1: Provide convincing evidence that the design, when implemented, will provide satisfactory mission support under real use scenarios and in the face of cyber-attacks.
  - This assurance case is supported by:
    - Rigorous logical arguments
    - Experimental evaluation
    - A detailed executable model of the design
- Phase 2: Use models to guide testing of implementation in increase security test effectivness
  - Test system aspects that are most important to overall system security

System Requirement: Design, Implement, and Validate a Publish and Subscribe Mechanism that …

- Provides 100% of critical functionality when under sustained attack by a "Class-A" red team with 3 months of planning
- Detects 95% of large scale attacks within 10 mins. of attack initiation and 99% of attacks within 4 hours with less than 1% false alarm rate
- Displays meaningful attack state alarms. Prevent 95% of attacks from achieving attacker objectives for 12 hours
- Reduces low-level alerts by a factor of 1000 and display meaningful attack state alarms.
- Shows survivability versus cost/performance trade-offs

Phase 1: Integrated Survivability Validation Procedure

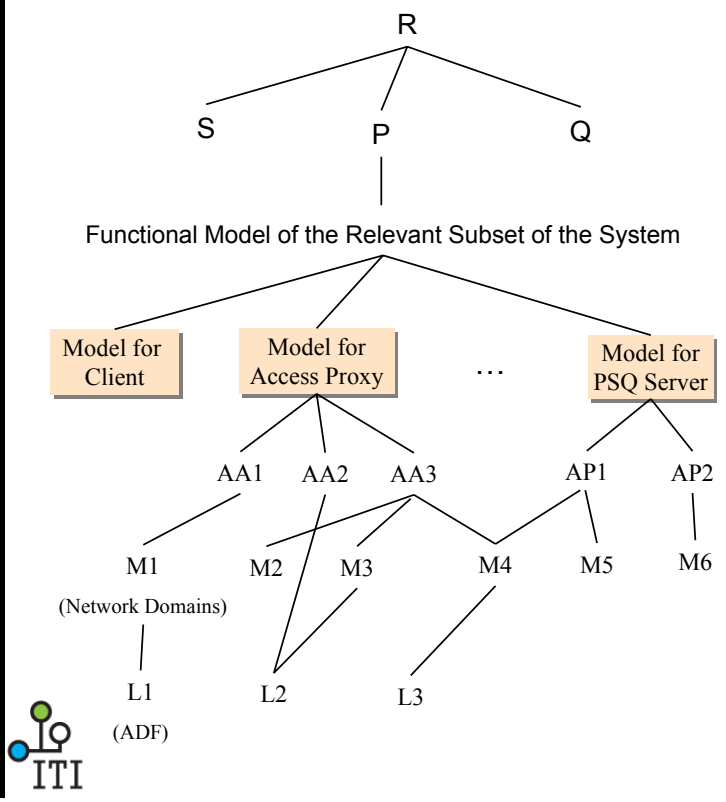

Requirement Decomposition

Functional Model of the System (Probabilistic or Logical)

Assumptions

Supporting Logical Arguments and Experimentation

Integrated Survivability Validation Procedure

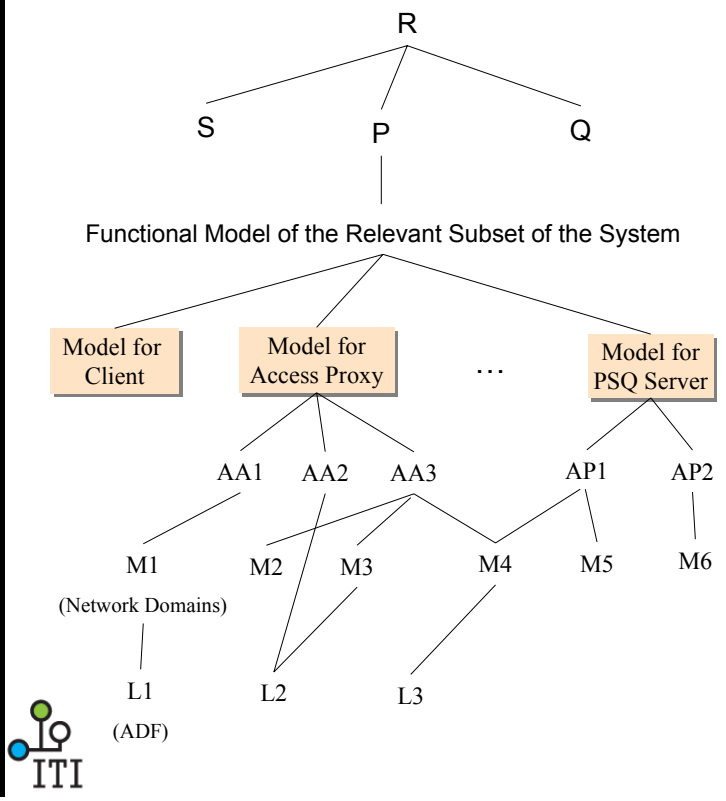

Steps

1. A precise statement of the requirements

2. High-level functional model description:
   a) Data and alerts flows for the processes related to the requirements,
   b) Assumed attacks and attack effects [Threat/vulner-ability analysis; whiteboarding]

Integrated Survivability Validation Procedure

Steps

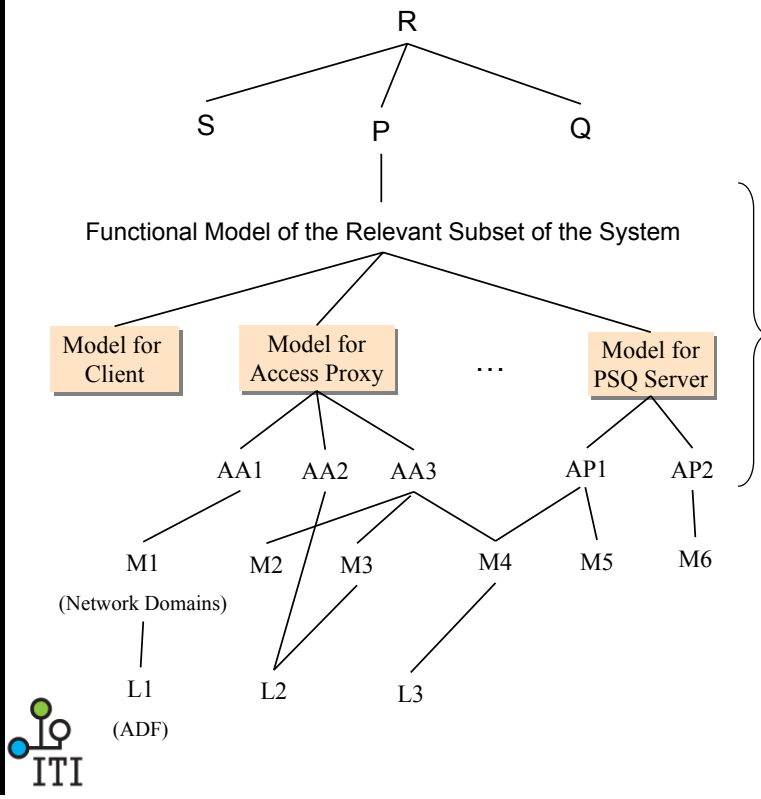

3. Detailed descriptions of model component behaviors representing 2a and 2b, along with statements of underlying assumptions made for each component. [Probabilistic modeling or logical argumentation, depending on requirement]

Integrated Survivability Validation Procedure

Steps

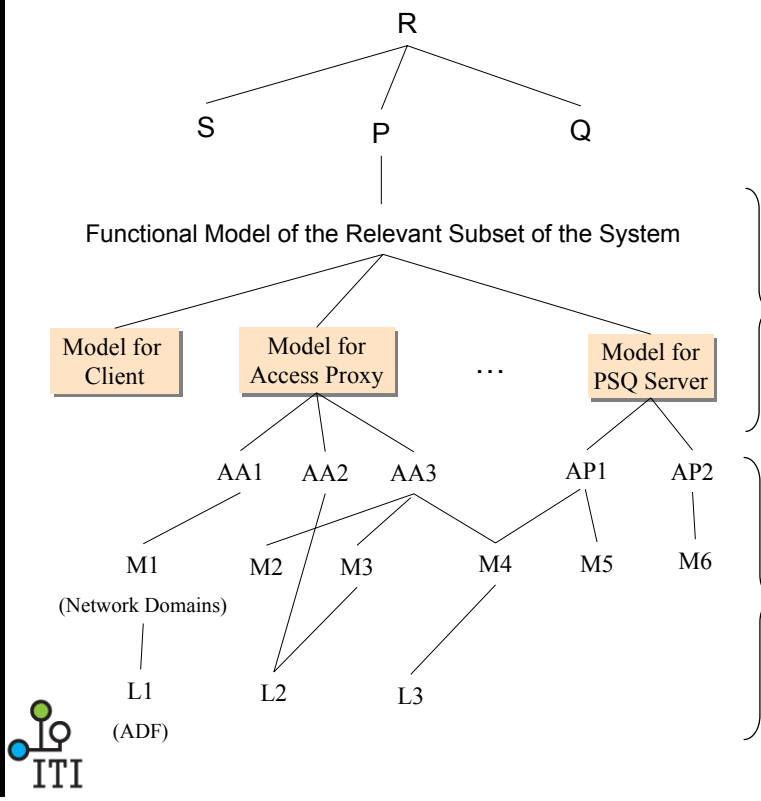

4. Construct executable functional model [Probabilistic modeling, if model constructed in 3 is probabilistic]

In Parallel

5. a) Verification of the modeling assumptions of Step 3 [Logical argumentation] and,
b) where possible, justification of model parameter values chosen in Step 4. [Experimentation]

Integrated Survivability Validation Procedure

Steps

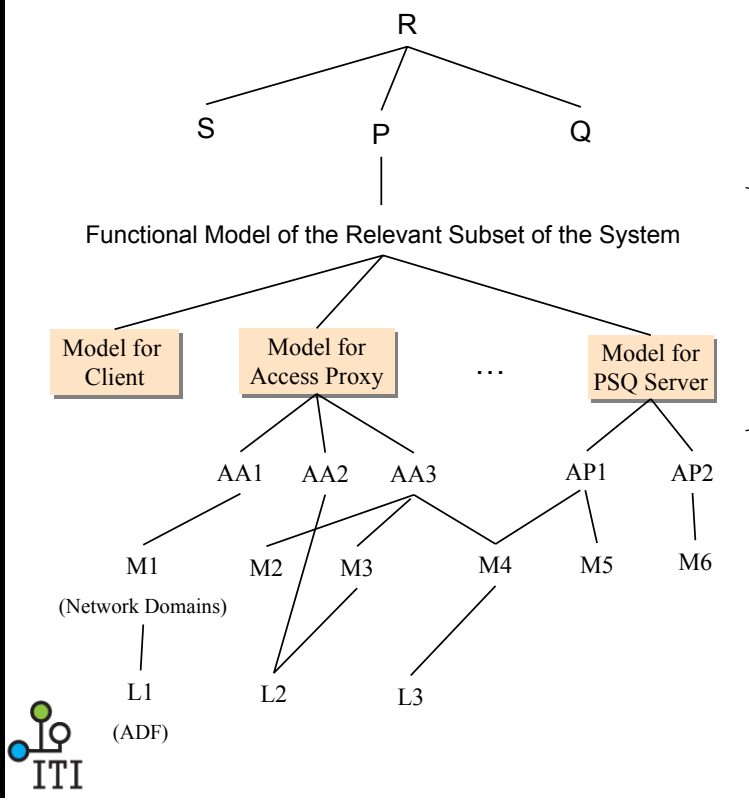

6. Run the executable model for the measures that correspond to the requirements of Step 1. [Probabilistic modeling]

Integrated Survivability Validation Procedure

Steps

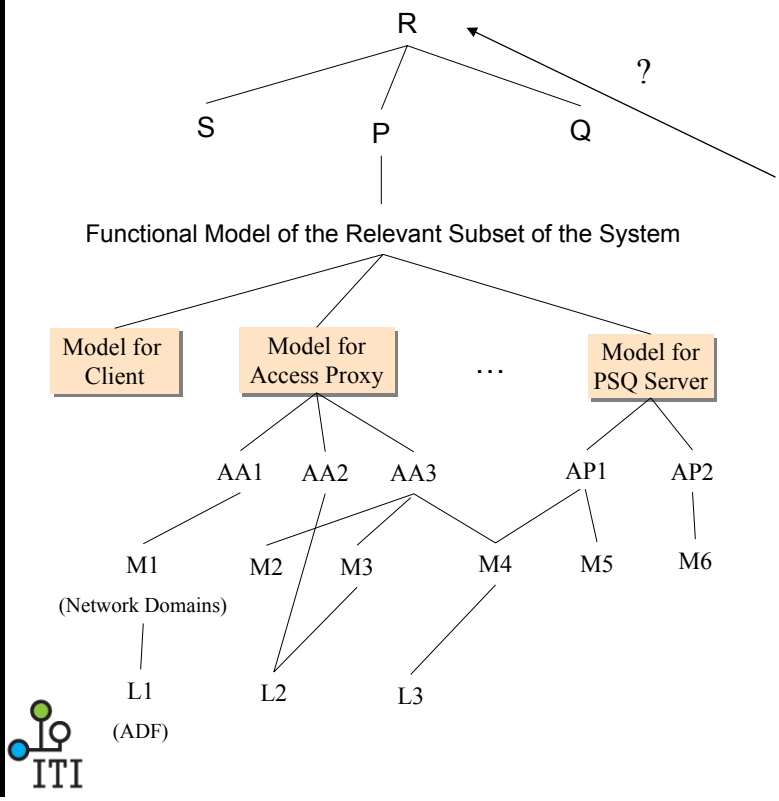

7. Comparison of results obtained in Step 6, noting in particular the configurations and parameter values for which the requirements of Step 1 are satisfied.

Note that if the requirement being addressed is not quantitative, steps 4 and 6 are skipped.

Argument Graph for the Phase 1 Design

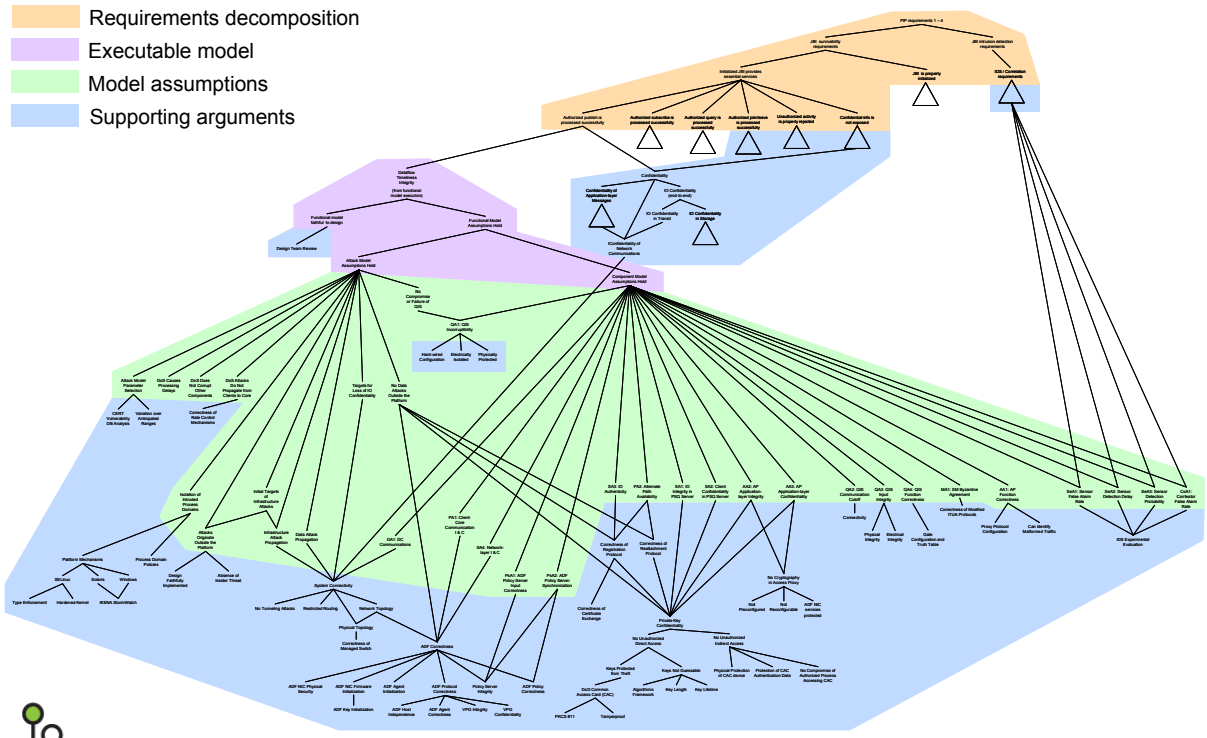

Attack Model Description

- Consider **effects** of attacks, not attacks themselves
- Attack propagation
  - MTTD: mean time to **discovery** of a vulnerability
  - MTTE: mean time to **exploitation** of a vulnerability
- 3 types of vulnerabilities:
  - **Infrastructure-Level Vulnerabilities** → attacks in depth
    - OS vulnerability
    - Non-JBI-specific application-level vulnerability
    - p_{common} : common-mode failure
  - **Data-Level Vulnerabilities** → attacks in breadth
    - Using the application data of JBI software
  - **Across process domains**
    - flaw in protection domains

Attack Effects

- **Compromise**
  - Launching pad for further attacks
  - Malicious behavior
- **Crash**
  - Attack propagation stopped
- **Distinction between OSes with and without protection domains**

Attack Response

- **Intrusion Detection**
  - p_{detect}=0 if the sensors are compromised
  - p_{detect} > 0 otherwise.
- **Attack Responses**
  - Restart Processes
  - Secure Reboot
  - Permanent Isolation

Infrastructure Attacks Example

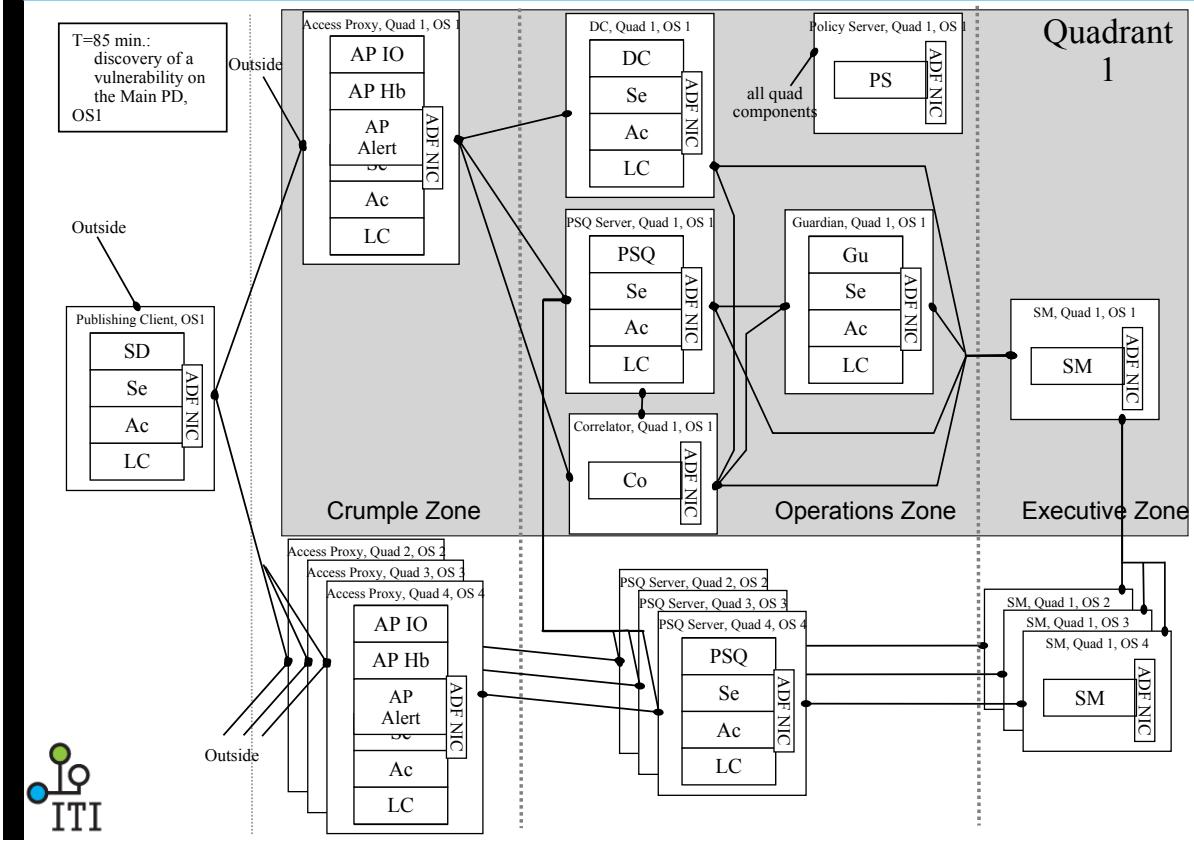

Construct Executable Functional Model

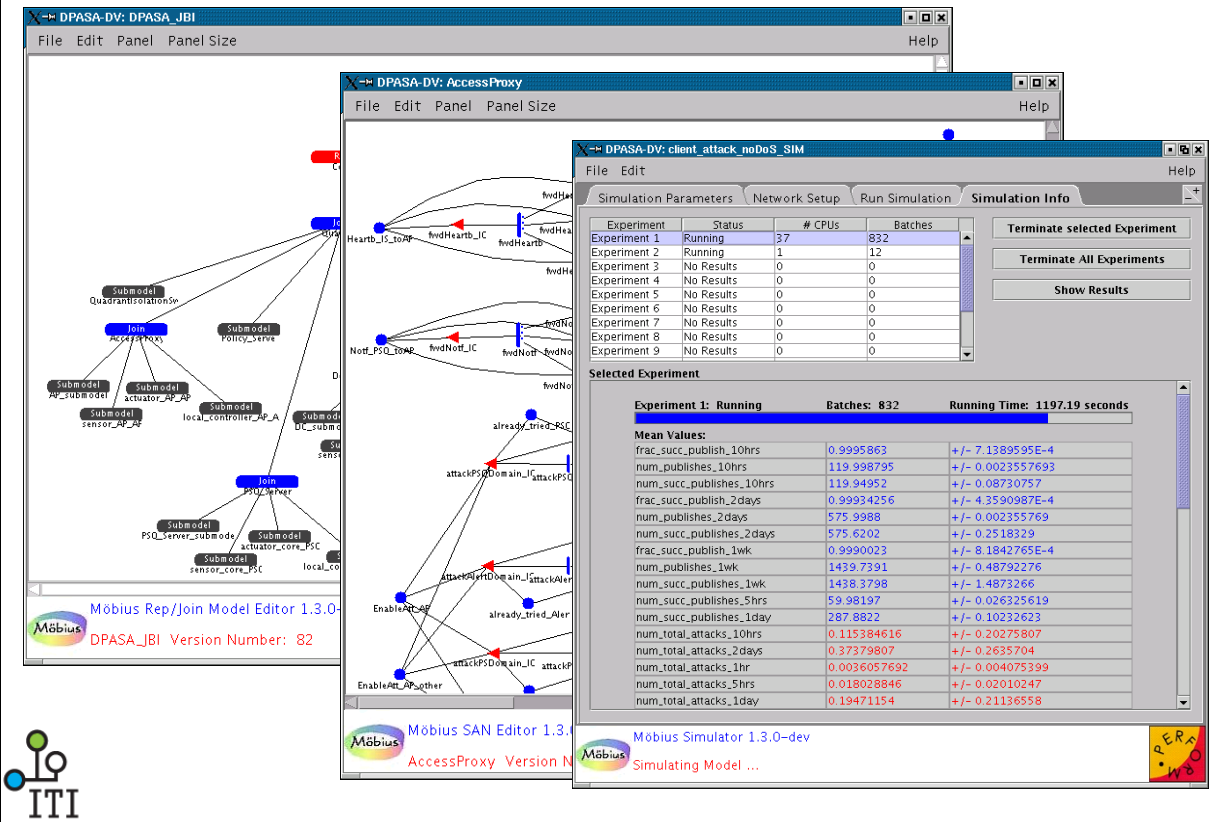

Vulnerability Discovery Rate Study

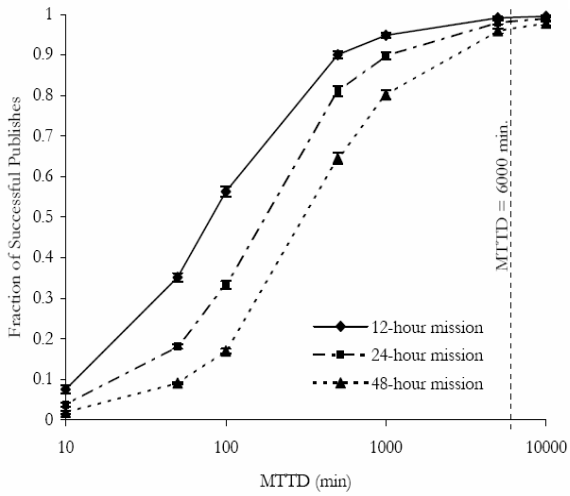

Fraction of successful publishes versus MTTD

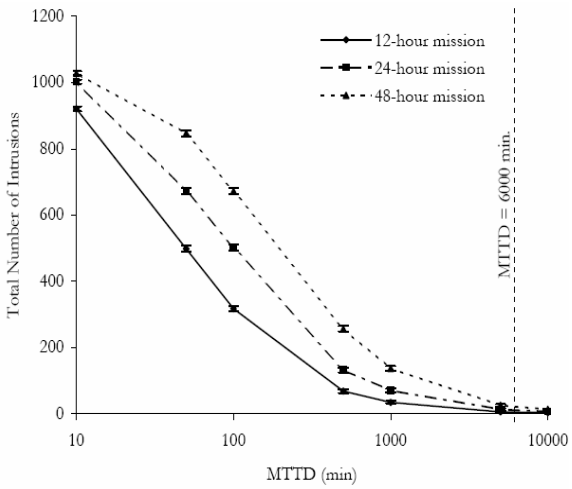

Number of successful intrusions versus MTTD

ITI

Varying the number of OS and OS w/ process domains

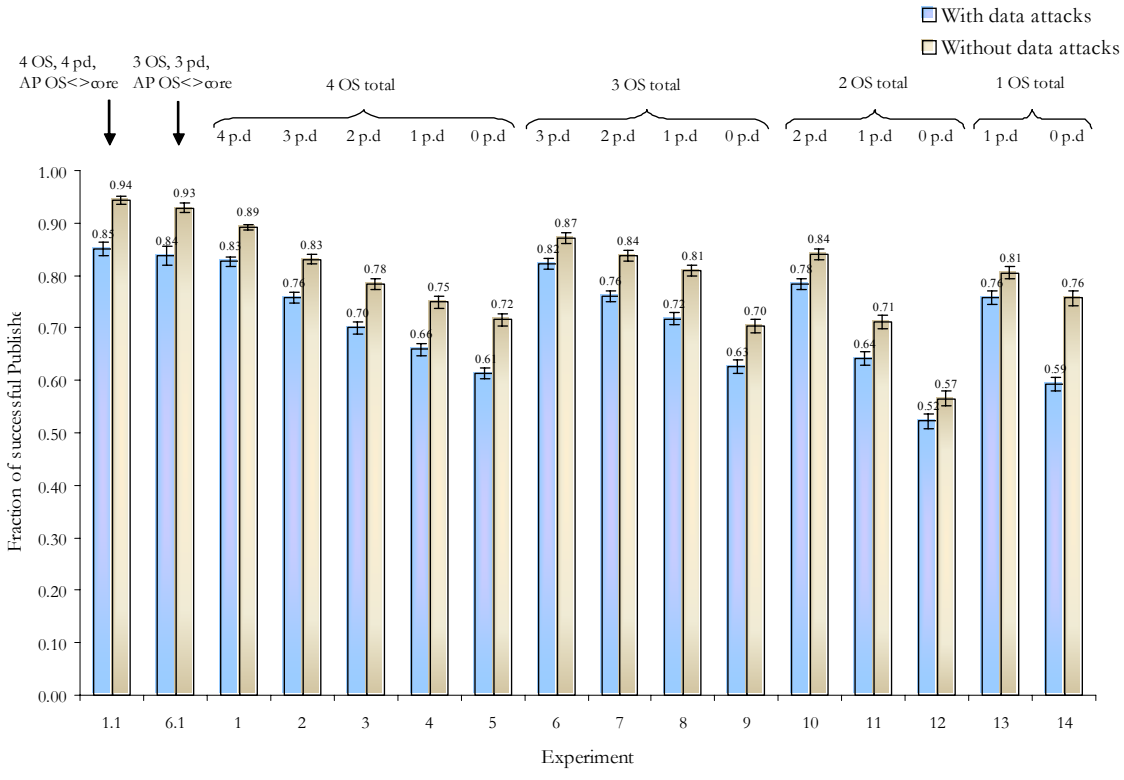

Phase 2: Improving (and Validating) the Implementation

Objectives:

- Improve the system's survivability
- Conduct specific system-level validation tasks
- Address all of the system-level concepts and mechanisms that may contribute to improvement, e.g., protocols and application scenarios

Main Idea:

- Think like an attacker
  - Examine whether a given attacker goal can be achieved
  - If so, alter the implementation so as to preclude such achievement

Procedure:

- Top-down, beginning with a specific high-level attacker goal
- Critical steps of the high-level attack tree are elaborated further as sub-trees, down to a level that admits adversarial testing.

Attacker Goals

- We considered the following attacker goals:

G1: Prevent client publish
G2: Prevent IO delivery to client (Subscription)
G3: Prevent a successful query operation
G4: Prevent a successful client registration
G5: Defeat confidentiality of IO data
G6: Modify IO data
G7: Modify data in repository

G5: Defeat Confidentiality of IO Data

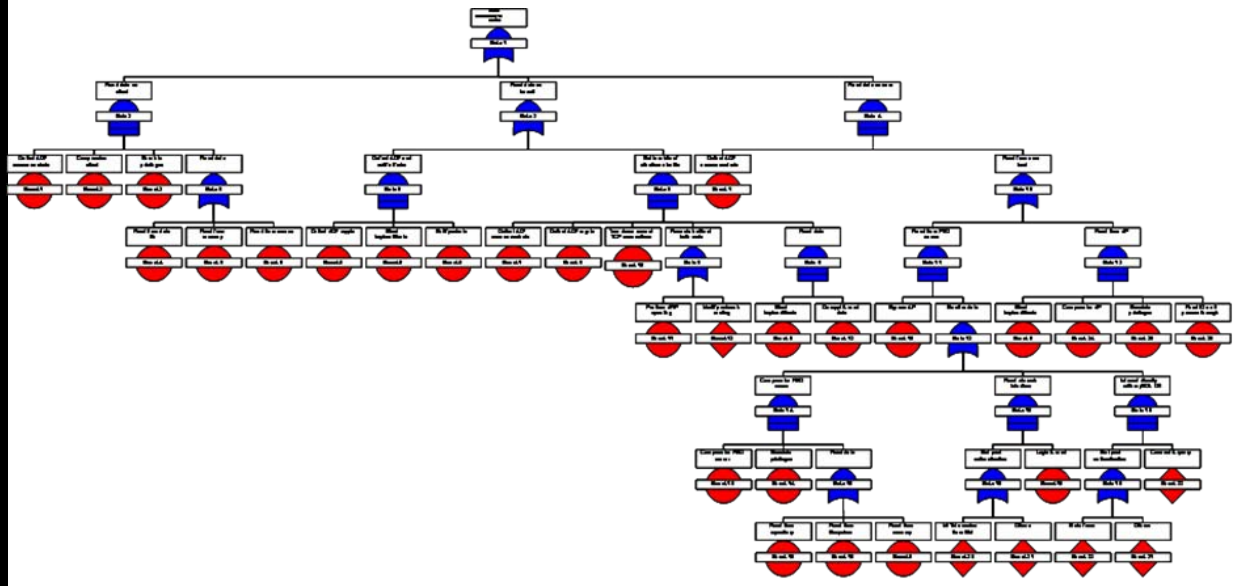

ITI

G5: Attack Steps/Minimal Attacks

Attack Step #	Type	Attack Step Description		Minimal Attack Sets
1 (3)	BASIC	Defeat ADF access control		7 , 8 , 9
2	BASIC	Compromise client		5 , 3 , 2 , 1
3 (3)	UNDEVELOPED	Escalate privilege		4 , 3 , 2 , 1
4	BASIC	Read from data file		6 , 3 , 2 , 1
5 (2)	BASIC	Read from memory		16 , 21 , 19 , 1
6	BASIC	Read from screen		16 , 20 , 19 , 1
7 (2)	BASIC	Defeat ADF crypto		16 , 21 , 22 , 1
8 (3)	BASIC	Steal key/certificate		16 , 23 , 22 , 1
9 (2)	BASIC	Sniff packets		
10	UNDEVELOPED	Tear down current TCP connections		
11	BASIC	Perform ARP spoofing		
12	UNDEVELOPED	Modify network routing		
13	BASIC	Decrypt & read data		
15	BASIC	Compromised PSQ server		
16	BASIC	Bypass AP		
17	BASIC	Read from filesystem		
18	BASIC	Read from repository		
19	BASIC	Login & read		
20	UNDEVELOPED	MITM session from SM		
21 (2)	UNDEVELOPED	Others		
22	UNDEVELOPED	Connect & query		
23	UNDEVELOPED	Brute force		
24	BASIC	Compromise AP		
25	BASIC	Read IO as it passes through		

ITI

Summary of Attack Steps/Minimal Attacks

- For the seven high-level attack trees that were developed, there are
  - 524 attack steps (including repeats)
  - 114 different attack steps
- The number of different minimal attacks for each high-level goal (these are derived automatically from a goal's attack tree) are as follows.
  - G1: 54, G2: 43, G3: 36, G4: 52, G5: 8, G6: 12, G7: 11
- Total number of minimal attacks with respect to all goals: 216

ITI

Attack Steps Frequency of Occurrence

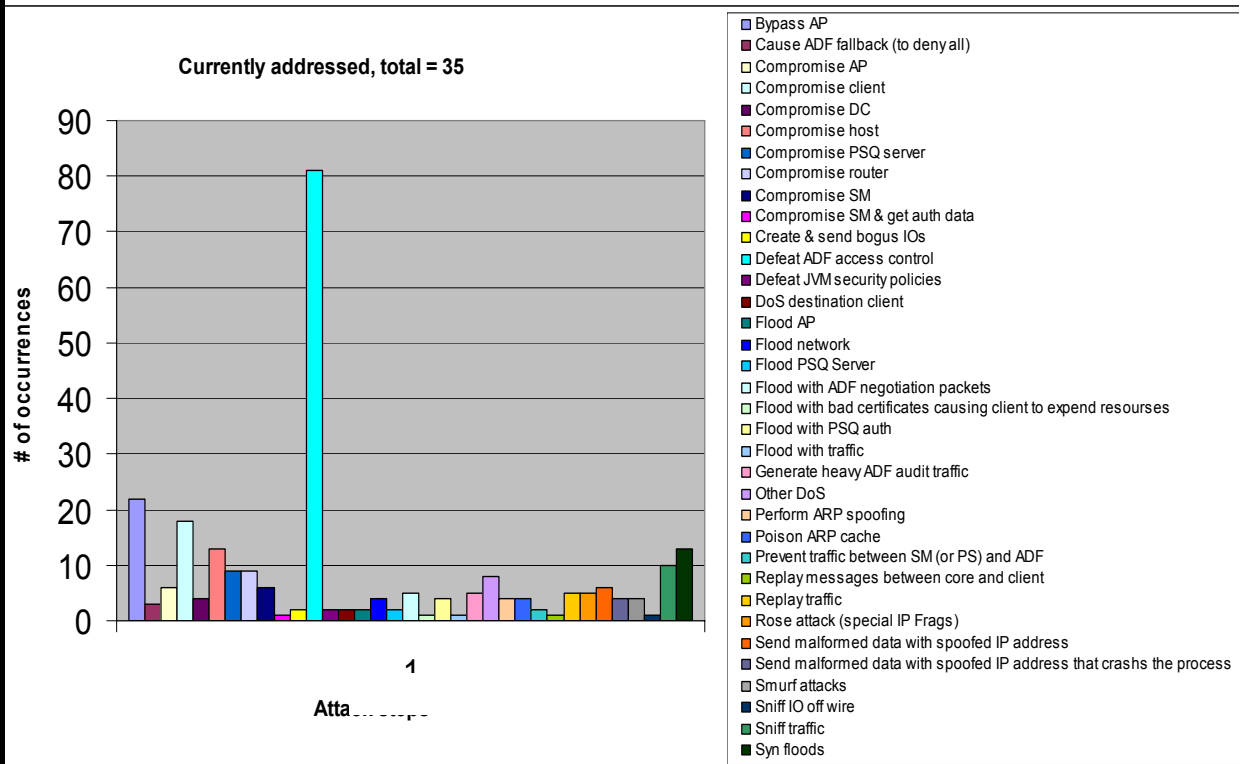

ITI

Example Attack Step Analysis: ADF DOS Attack

- Three Metrics were used to benchmark the ADF.
  - *Max. Throughput:* The fastest receive rate at which there is no packet loss
  - *Available Bandwidth:* The amount of data that can be transmitted in a fixed amount of time (when no flood in progress)
  - *Minimum Flood Rate:* The lowest rate of flood which leads to a successful denial of service attack.
- Floods cause packet loss, which in turn lowers bandwidth due to TCP congestion control. UDP will suffer high packet loss.
- Experimental Setup
  - Follows rfc2544 as much as possible
  - Max flood rate is ~44000 frames/sec = 22 Mbits/sec (for 64 Byte frames)

ITI

Minimum Flood Rate for Successful DoS on ADF NIC

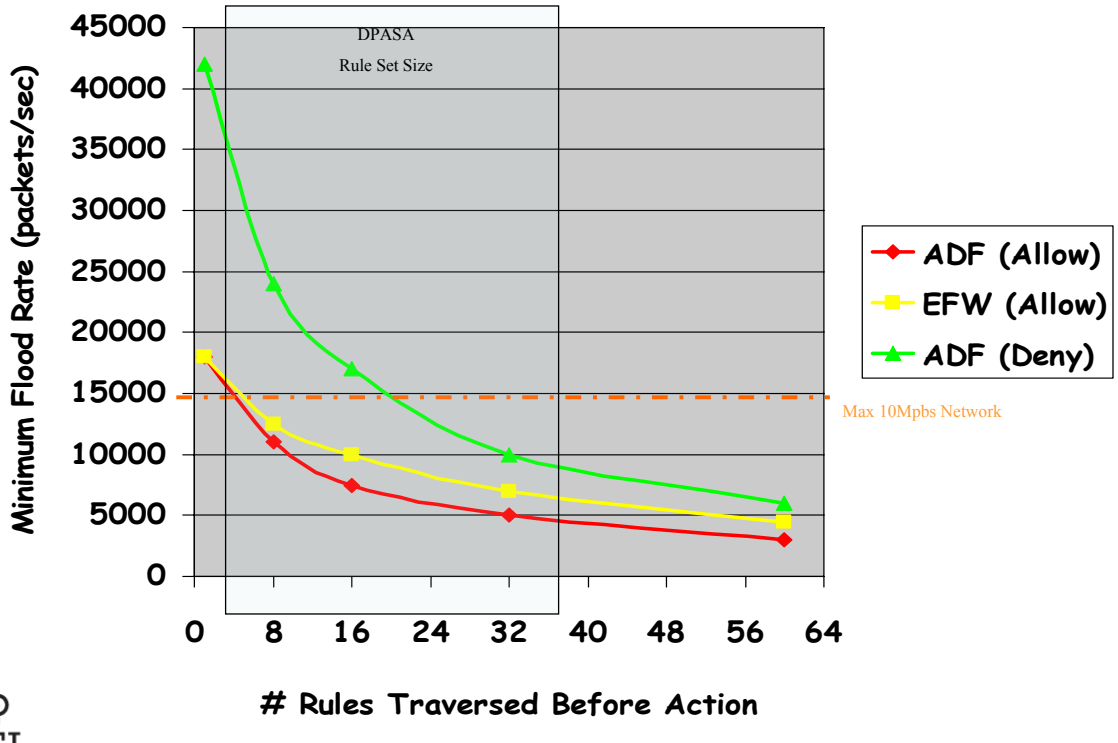

ITI

Conclusions

- How can the behavior of attackers be quantified?
  - By their effect, if system is intrusion tolerant
- How should security/survivability measures be specified?
  - In terms of the definition of "proper operation" for the system
- If relative measures are desired, can they be shown to be robust across a wide variety of situations?
  - Yes, through extensive simulation
- How accurately can absolute measures be estimated?
  - Unknown ???
- Can quantification aid in security testing?
  - Yes, through (advanced) attack tree analysis
- Can a notion of "coverage" be developed for security testing?
  - Unknown ???

Modelling and Evaluation of Largeness in Evolving Systems

Andrea Bondavalli

University of Firenze (here PISA)

Introduction

Systems **complexity** has always been a very critical issue and is becoming even worse in modern infrastructures and systems.

When modelling such systems, complexity of the resulting models depends on the

- dependability measures to be evaluated,
- the modelling level of detail, and
- the stochastic dependencies among the components.

State-space models are commonly used and require a very high number of states for the modelling **and complex and costly analytical techniques, or simulation for they solution**

The **large size** of models known as the **'state space explosion problem'** is one of the major difficulties in the dependability evaluation of real systems.

How to cope with Largeness

Much work done and progress made in addressing such problems at the model construction and model solution levels.

These are complementary and both are needed to generate and process detailed and large dependability models for the evaluation of the resilience of real life systems.

In the rest of the presentation we will illustrate

- Three main classes of structured techniques for a modular model construction.
- Model solution techniques.
- Specific methods developed to deal with such large and evolving systems taking as examples web, grid and mobile based systems

Model construction techniques

At model construction level, we can identify three approaches:

i) model composition; the system model is constructed in a bottom-up fashion. The models representing parts of the systems are built in isolation, thus having a limited view of the system context.

ii) system decomposition and model aggregation; it follows a top-down approach: starting from an overall view of the system context, the model for the overall system is decomposed in a set of simpler sub-models.

iii) the derivation of dependability models from high-level specifications (based on UML -Unified Modeling Language- or AADL -Architecture Analysis and Design Language- the overall model (e.g., a Markov chain or a Petri net), is built by transformation (usually semi-automatic) from such high level specification.

Composition approaches - Principles

The principle of the composition approach is
- to build complex models in a modular way through a composition of its sub-models
- then solved as a whole.

Most of the works belonging to this class define the rules to be used to construct and interconnect the sub-models
- exploiting the degree of dependency among subcomponents.

These dependencies are used to reduce the model complexity creating, smaller, equivalent representations.

Example

Stepwise refinement approach [Betous-Almeida & Kanoun 2004-a] following the system development refinement process.

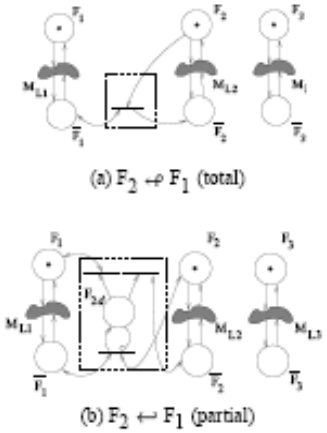

Figure 4. Functional dependencies

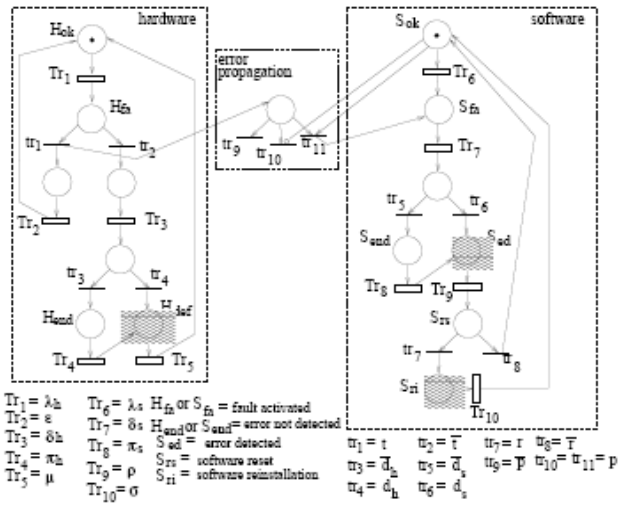

Figure 15. Structural model of a software and a hardware components

Decomposition / aggregation approaches - Principles

Most decomposition and aggregation methods are characterized by a hierarchical decomposition approach

Thus they try to avoid the generation of large models.

The overall model is decoupled in simpler and more tractable sub-models:

-sub-models are solved separately and

-the measures obtained from the solution of the sub-models are then aggregated to compute the overall measures.

An example:

Decomposition exploiting the hierarchy of control systems [Lollini *et al.* 2005-a], solution carried out in a bottom-up fashion (aggregation).

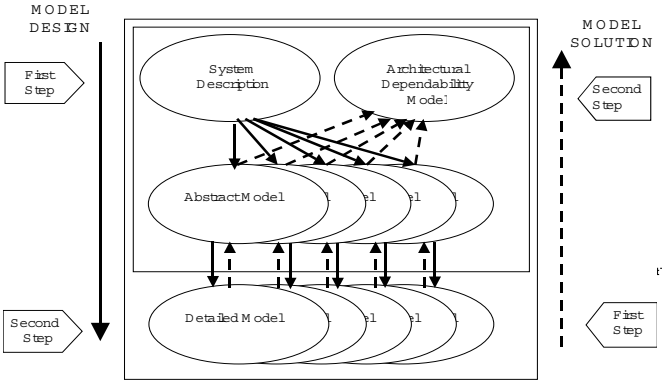

Detailed generic Model

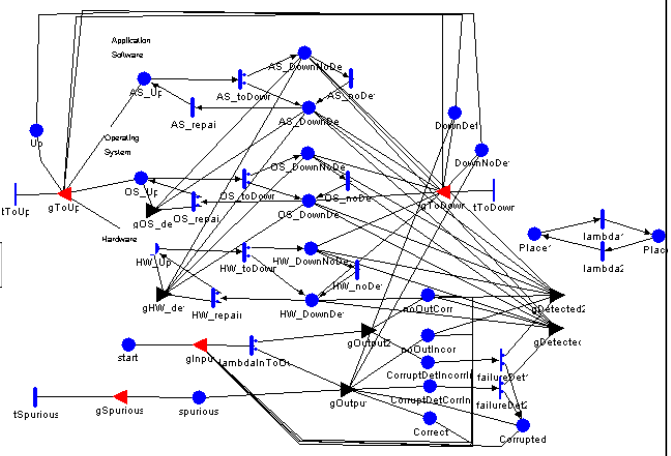

Derivation from high-level specification

Model-driven engineering are more and more used in industry (in particular UML and AADL)

As system designers use integrated set of methods → approaches have been developed for allowing the (semi-automatic) generation of dependability evaluation models from such model-driven engineering.

Research based on UML:

- The European project HIDE [Majzik & Bondavalli 1998, Bondavalli *et al.* 2001a, Majzik *et al.* 2003] automatic analysis defining several model transformations from structural and behavioural UML diagrams into GSPNs, DSPNs and SRNs.
- The issue of deriving automatically models from UML behavioural specifications, has also been addressed in [Bernardi 2003].
- synthesis of dynamic fault trees (DFT) from UML system models [Pai & Dugan 2002].

AADL has more recently received some interest:

- A stepwise approach for the description of complex dependability models from AADL [Rugina et al. 2006].

Solution approaches

Two main approaches for dealing with largeness at solution time

- **largeness avoidance techniques** that try to reduce the size of the generated models
- **largeness tolerance techniques** which make use of space and time efficient algorithms to reduce the storage requirements of the state space and the generator matrix and to optimize the state space exploration, generation and analysis.

It is important to note that largeness avoidance and largeness tolerance techniques are complementary

Both are needed, at model construction and model solution levels each bringing its contribution.

Large and evolving systems

Systems are evolving and becoming more complex and large.

They are also more and more closely interconnected and show increasingly complex interactions.

All this is demanding a continuing evolution and improvement of the modelling and evaluation capabilities in order to quantify their dependability characteristics.

Among types of systems that present these challenges we considered

- Dependability modelling of Web-based systems and services
- QoS analysis of Mobile Telephone Systems
- Service Provisioning and Grid Systems

Dependability modelling of Web-based systems and services

The dependability of the **delivered services as perceived by the users** is a key issue for Internet applications and Web Services

When Internet is used for money critical applications (online banking, stock trading, reservation processing and shopping) Availability (wrt. Accidental & Malicious faults) is critical.

- Many measurement-based efforts for the evaluation of the of web hosts [Oppenheimer & Patterson 2002, Kalyanakrishnam et al. 1999], less emphasis put on modelling.
- A multi-level approach for modelling the user perceived availability of internet applications considering 4 abstraction levels modeled with various techniques, [Kaâniche *et al.* 2003-a].
- Detailed analytical performability models to analyze the availability of web services implemented on cluster architectures. [Martinello et al. 2005]
- Dependability modelling of web-based systems and services performed considering a business model workflow [Gönczy et al. 2006]:

QoS analysis of Mobile Telephone Systems

Information Society Technologies

SIXTH FRAMEWORK PROGRAMME

Telephone Systems are getting more and more business critical and complex showing strong interactions with an integrated Information and telecommunication Infrastructure.

- Analysis of GPRS by providing a modelling approach to understand the effects of outage periods on the service provision [Porcarelli et. al. 2002, Porcarelli et. al. 2003].
- Congestion analysis of GPRS infrastructures consisting of a number of partially overlapping cells [Lollini *et. al.* 2005-b], using QoS indicators as a measure of the service availability perceived by users.
- A general approach [Lollini *et. al.* 2006] applicable to cellular systems, including GSM, GPRS and UMTS networks. It enhances the modularity, reusability, scalability and the maintenance of the overall model.

Service Provisioning and Grid Systems

Information Society Technologies

SIXTH FRAMEWORK PROGRAMME

Various novel IT business models depend on adaptive infrastructure mechanisms to share resources, create distributed and collaborative applications, and manage and maintain systems and applications.

Such platforms, (known as grid computing, service provisioning, utility computing, on demand computing) pose various new challenges on evaluation methods and techniques.

- In [Jarvis et al. 2004] various new challenges with respect to the performability evaluation of such systems are addressed
- In [Palmer & Mitrani 2005], theoretically optimal policies to allocate resources to customers are computed, and compared with a newly proposed heuristic validated and tuned using the experimental system described in [Fisher et al. 2004].
- A methodological approach [Machiraju et al. 2002] to systematically introduce metrics for the business' operation and managers. It relies on the concept of 'Quality of Business' [van Moorsel, 2002], and is implemented based on contracts and/or service level agreements (SLAs) [Molina et al. 2005].

Conclusions

- The increasing scale and complexity of modern-day computing systems continues to demand good techniques for the construction and solution of large quantitative models.
- In addition, these large, dynamic and evolving systems pose some **new challenges** that the ReSIST partners aim to address.
- Evaluation methods must deal with metrics at an increasingly high level of abstraction, to express the impact of the computing infrastructure on an enterprise business.
- Of increased significance is also the need of quantitative evaluation methods to support the effective use of adaptation mechanisms prevalent in modern-day systems.

Towards attack modelling thanks to honeypot data processing

Marc Dacier

Institut Eurécom

Sophia Antipolis, France

dacier@eurecom.fr

Overview

- Introduction
- *State of Knowledge*
- *Contributions of ReSIST Partners*
- *Conclusions*

Threats?

- *Fact:* New vulnerabilities discovered every day, new widespread attacks reported in the media.

- Questions:
  - Are these vulnerabilities actually exploited?
  - What are the “right” fault assumptions models that one should use to build intrusion tolerant systems?

Dahu: definition

source: http://www.vidonne.com/html/dahu-reignier.html

“The Dahu is an extremely shy animal living in the Alps of France and Switzerland.[…] It has adapted to its steep environment by having legs shorter on the uphill side and longer on the downhill side […]
“

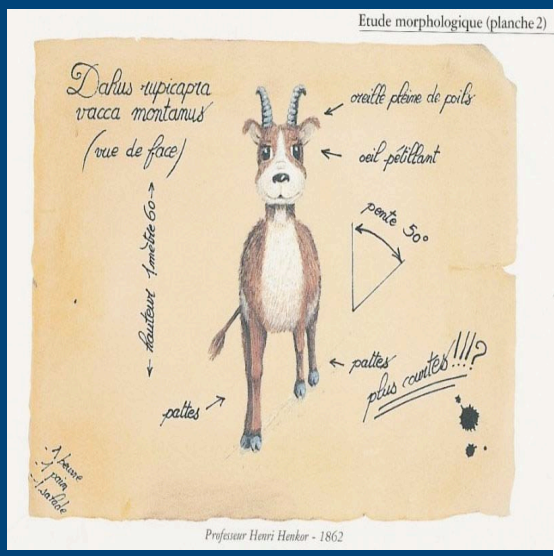

“The Dahu, An endangered Alpine species”, *Science*, 2568, November 1996, pp.112:

Food for thoughts …

- *Dahus* are rare, bizarre, stimulating from an intellectual point of view but ...
- Does it justify the existence of *Dahusian research*?
- What about *Dahusian research* in security assessment?

Overview

- Introduction
- *State of Knowledge*
- *Contributions of ReSIST Partners*
- *Conclusions*

The basics

- « A Honeypot is an information system resource whose value lies in unauthorized or illicit use of that resource »

L. Spitzner, *Honeypots: tracking hackers*, Addislon Wesley, 2002

The basics (ctd.)

- Low interaction honeypots:
  - emulate the existence of a potential target,
  - At various abstraction levels (network, OS, application)
- High interaction honeypots:
  - Use a real system as a potential target
  - Must be kept under close scrutiny.

Internet Telescopes

- Internet Telescopes observe empty address spaces:
  - CAIDA Telescope,
  - IMS,
  - iSink,
  - Minos,
  - Team Cymru,
  - Honeytank,
  - IUCC/IDC Internet Telescope (Israel),
  - Etc...
- The Honeynet Alliance promotes the use of high interaction honeypots.

Problems with current solutions

- **False positives**
  - **It may be difficult to discriminate true attacks from erroneous, yet legitimate behaviours, in data collected in real networks.**
- **Privacy**
  - **Data sets may contain private information (eg IP addresses, passwords, etc.). Anonymisation removes semantic and is therefore not always usable.**
- **Liability**
  - **Not stopping an ongoing attack may harm third parties. Major issue for high interaction honeypot.**

Problems with current solutions (ctd.)

- **Bias**
  - **Things may be different here and there.**
  - **Malicious users dislike to be observed and will avoid visiting known observation points (eg .mil, major corporate networks, etc..)**
- **Amount of data**
  - **Having access to a large amount of data is good**
  - **Having access to a rich amount of data is better.**
  - **Having access to a rich amount of complete and comparable data is even better!**

Summary

- What we need is:
  - an environment to collect unbiased, rich, complete and comparable data about attacks without facing liability or privacy issues.
- To do so, we have deployed:
  - the very same low interaction honeypots in a large number of diverse locations using each time a very limited amount of IP addresses. We collect all packets sent to or from these machines, including payload.

Overview

- Introduction
- *State of Knowledge*
- *Contributions of ReSIST Partners*
- *Conclusions*

Collaborative approach

- Leurré.com framework used as a common umbrella to carry out joint research in this thema.
- Some partners bring also on the table the expertise gained with their own proprietary dataset (eg. IBM with its internal Billy Goat project).

50 partners in 30 countries covering the 5 continents

Experimental Set Up
(based on honeyd)

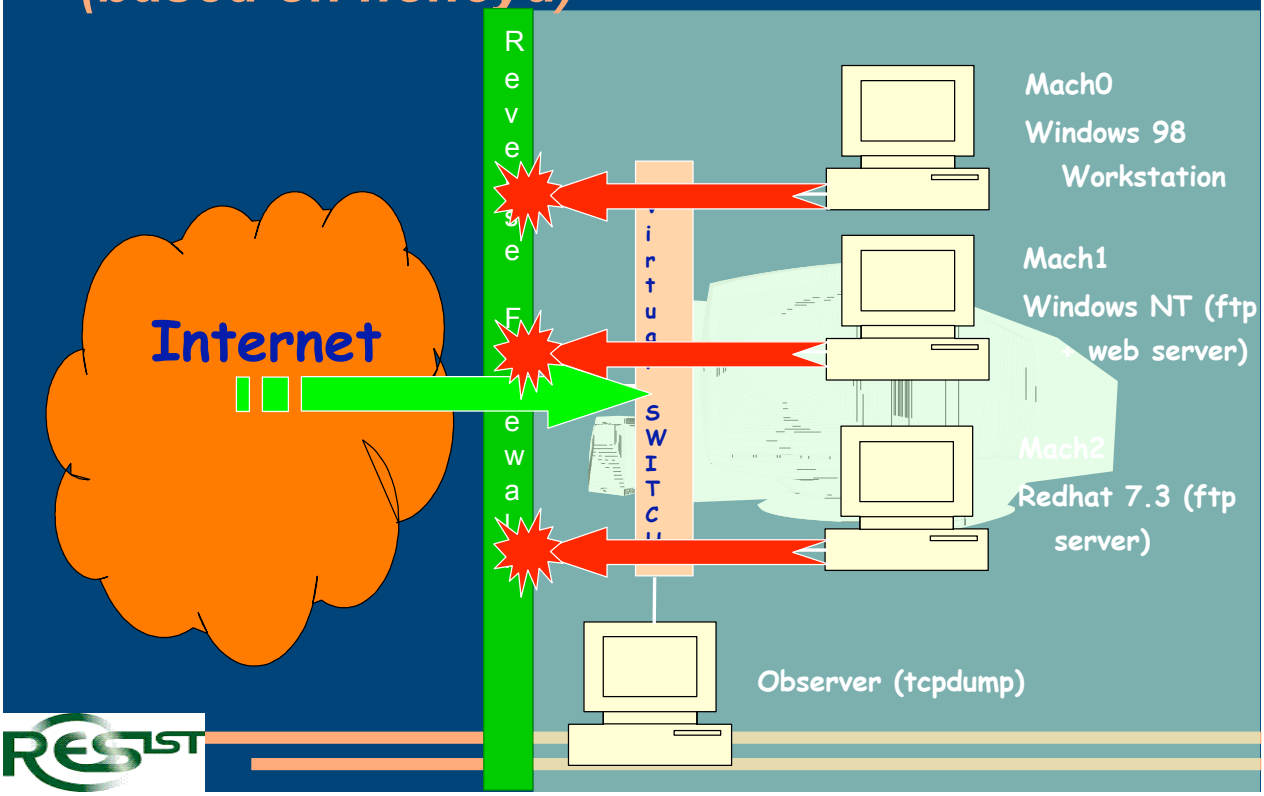

Win-Win Partnership

- The interested partner provides …
  - One old PC (pentiumII, 128M RAM, 233 MHz…),
  - 4 routable IP addresses,
- The project offers …
  - Installation CD Rom
  - Remote logs collection and integrity check.
  - Access to the whole SQL database by means of a secure GUI and a wiki (over https).

D12 - Appendices

- [Alata et al. 2006] E. Alata, V. Nicomette, M. Kaaniche and M. Dacier, “Lessons learned from the deployment of a high-interaction honeypot”, Proc. Sixth European Dependable Computing Conference (EDCC-6), Coimbra, Portugal, October 18-20, 2006

- [Kaâniche et al. 2006] M. Kaâniche, E. Alata, V. Nicomette, Y.Deswarte, M. Dacier, “Empirical analysis and statistical modelling of attack processes based on honeypots”, Proc. of WEEDS 2006 - workshop on empirical evaluation of dependability and security, Philadelphia (USA), June 25 - 28, 2006.

[Alata et al. 2006]

- High interaction honeypots are not that rapidly detected.
- They help in identifying groups of attackers and their strategies.
- They are complementary to low interaction ones
- Very difficult to use to collect long term datasets.

[Kaâniche et al. 2006]

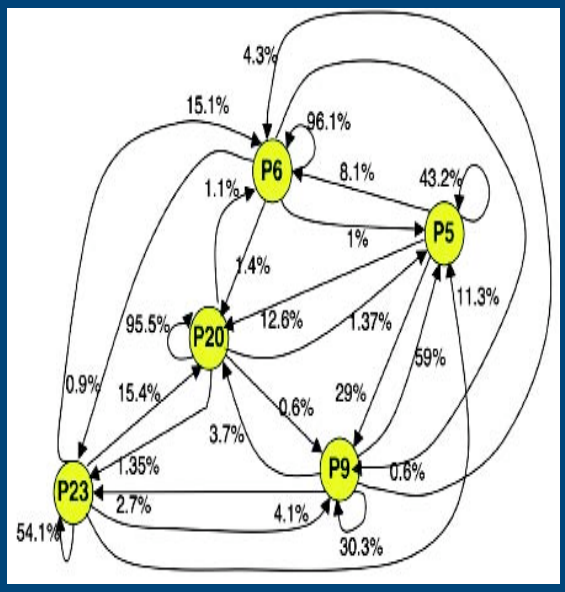

- Propagation graphs open the way to predictive models for *some attacks*

[Kaâniche et al. 2006]

- Patterns of attacks common to several platforms open the way to predictive models for *some platforms* (20/12/06 - 31/1/07)

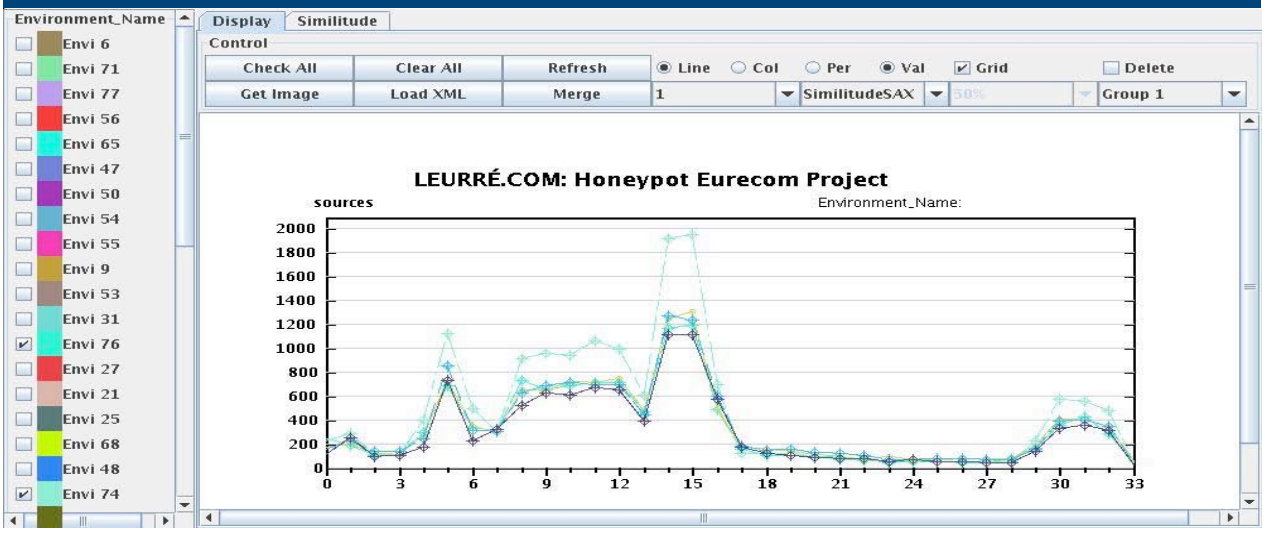

Overview

- Introduction
- *State of Knowledge*
- *Contributions of ReSIST Partners*
- *Conclusions*

Conclusions

- First results demonstrate the usefulness of such datasets with respect to the proposed objectives.
- Honeypots with higher degree of interaction would be welcome.
- Models must be formalized and validated.

Scalable Verification of Systems with Cryptography

Birgit Pfitzmann (IBM Research, Zurich)
Joint work mainly with Michael Backes (Univ. Saarbrücken) and Michael Waidner (IBM Research & SWG)

Building Secure Systems

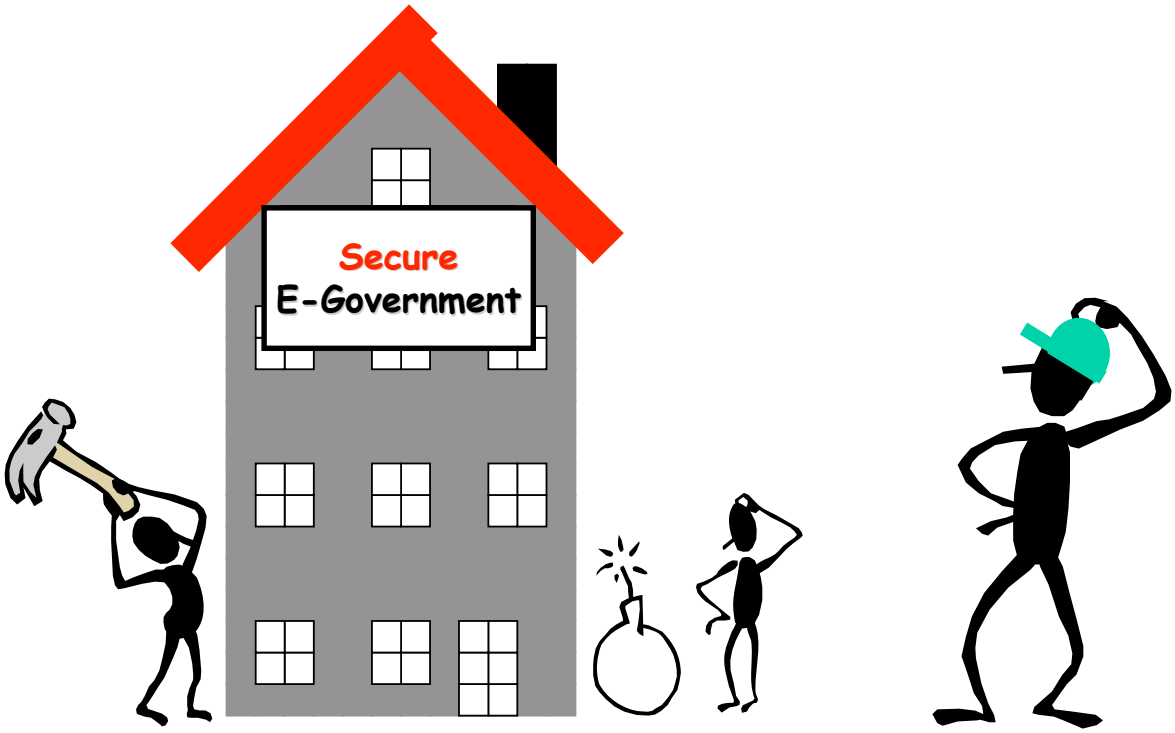

Cryptography: The Details

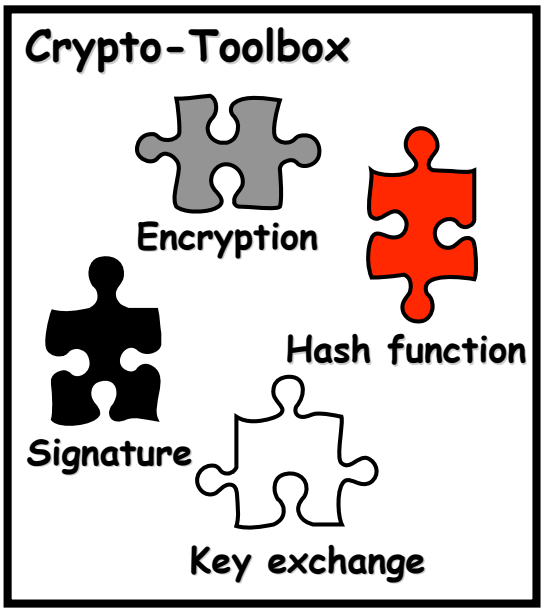

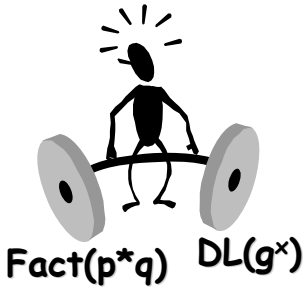

Cryptography: The Details

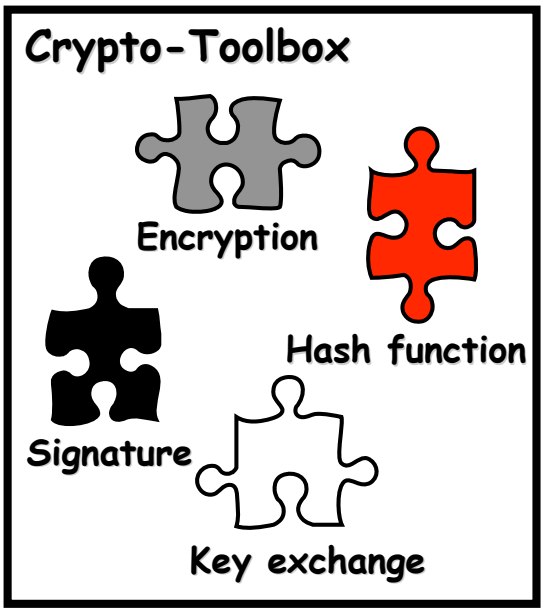

Prior Automated Crypto Protocol Proofs: The Big Picture

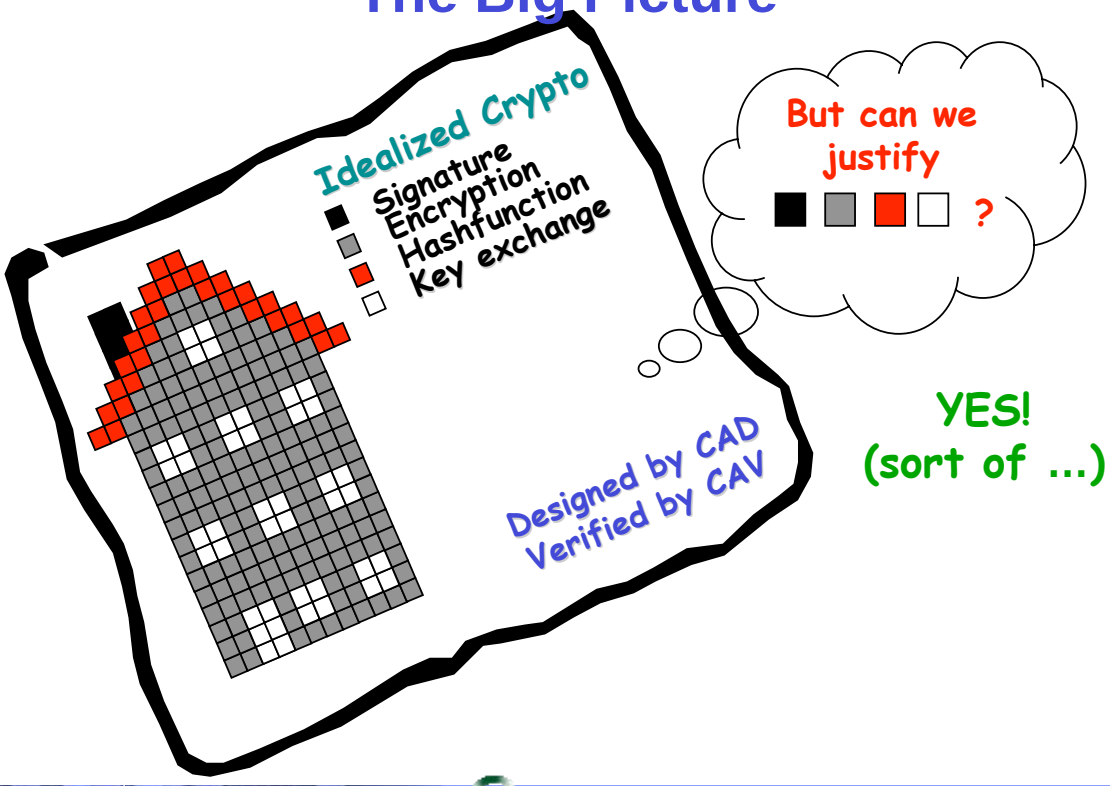

E.g.: Secure Channels like SSL
(with mutual authentication)

- **If you use them in a larger system, what would you assume about them, or how would you model them?**
- **E.g., as "ideal secure channel"**

  - **E.g., as a primitive in π-calculus etc.**

Secure Channels, ctd.

- How correct is this compared with actual SSL?
- Not bad, but not quite correct:
  - Computational assumptions and error probabilities from crypto
  - Message length and traffic pattern leak
  - No availability

Always very similar
⇒ make part of semantics ("fulfillment" relation)

Special
⇒ extend specification

Rather general
⇒ can just be in asynchronous model

Reactive Simulatability (RSIM)

Here "General RSIM" variant

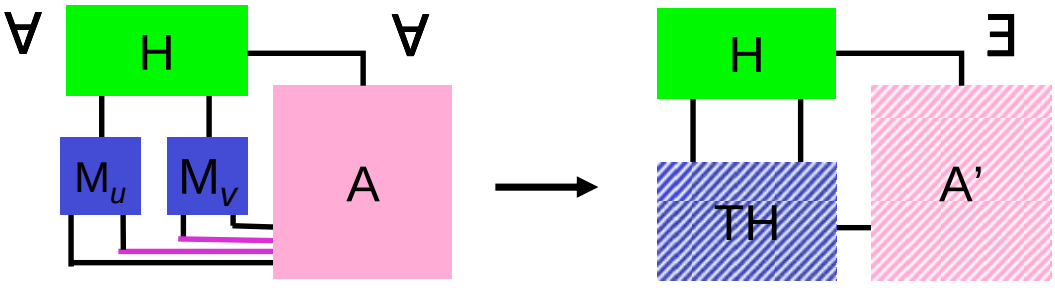

Real system **Ideal system**

$$view_{real}(H) \approx view_{ideal}(H)$$

Indistinguishability of random variables

RSIM in Overall Design Process

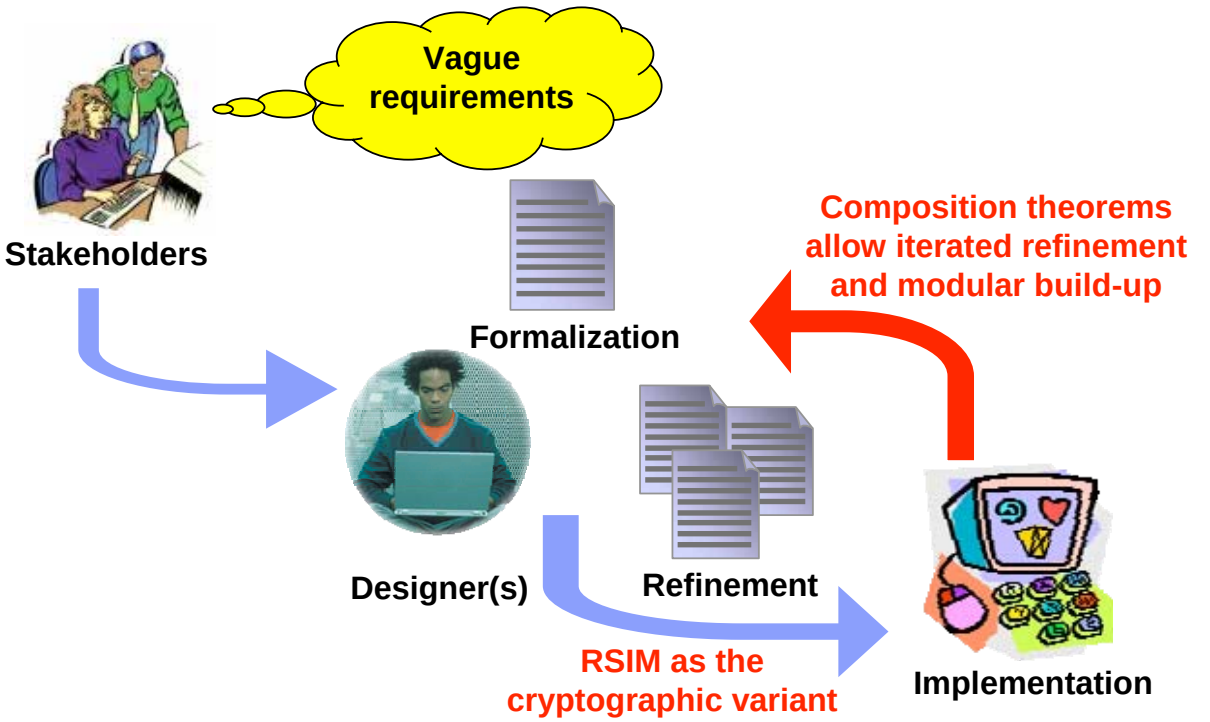

Treating Properties Cryptographically

- Integrity requirements
- No information flow between certain system parts
- Secrecy of specific data

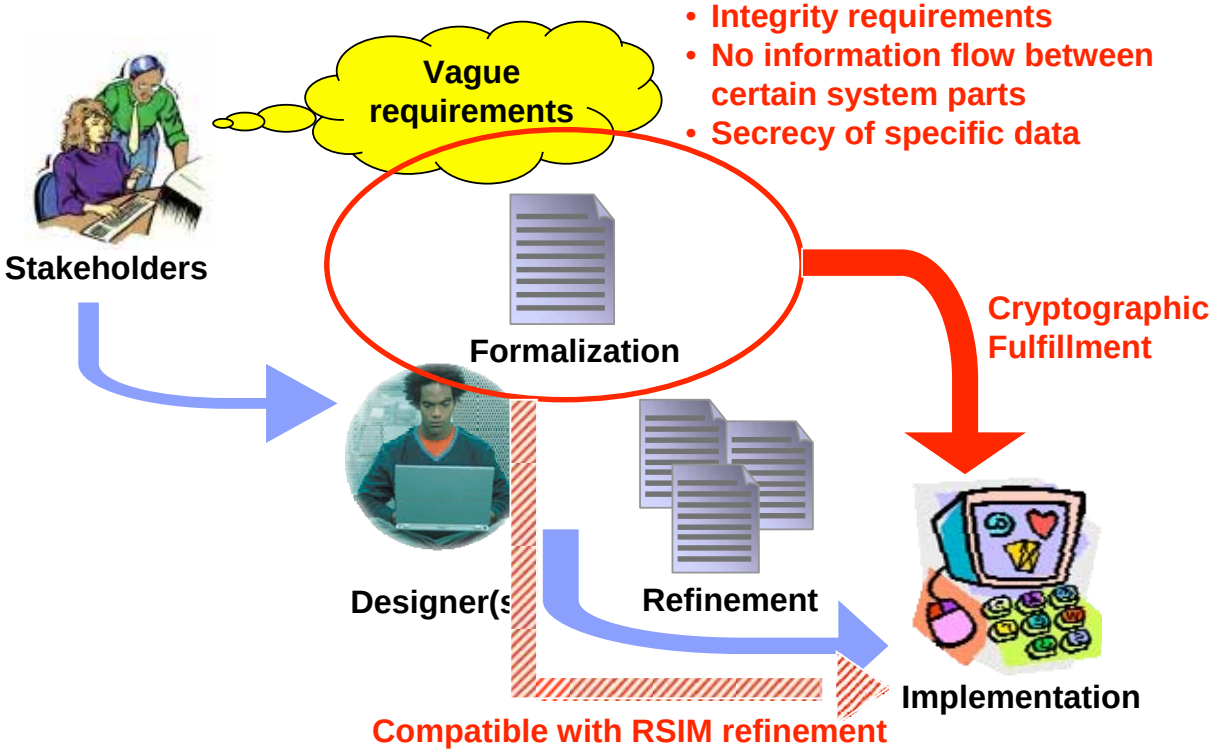

Recent Work

- **Extended prior results for “Dolev-Yao models” – specific term-algebra abstractions widely used in verification community**
- **Impossibility results for certain Dolev-Yao model variants**
- **BPW-Dolev-Yao model in Isabelle/HOL (with Ch. Sprenger and D. Basin)**
- **Attempt to apply to real-world Web Services**

Prototype Knowledge Base: an on-line information service in dependability and security

Hugh Glaser
Electronics & Computer Science
University of Southampton

Budapest, 22nd. March 2007

Information Society Technologies

SIXTH FRAMEWORK PROGRAMME

With

- Ian Millard
- Afraz Jaffri
- Benedicto Rodriguez
- ReSIST Partners
  - esp. Brian Randell

Background: Semantic Web Challenge 2003 Winner

- CS AKTive Space
  - Gather data
  - UK People, projects, publications
    - Research funding
    - Top Universities
  - Geographical presentation
- AKT Project (www.aktors.org)

The Challenges

- Scientific Intelligence
  - Who is doing what where?
  - What impact are they having?
- Integrating resources
  - CORDIS, Institutional DBs and web sites, ePrints, NSF, CiteSeer, RISKS list, ISO LoCodes…
- Information: distributed and heterogeneous
  - Not under own control
  - Not in a common format
  - Not where you expect it
- Presenting to users & agents

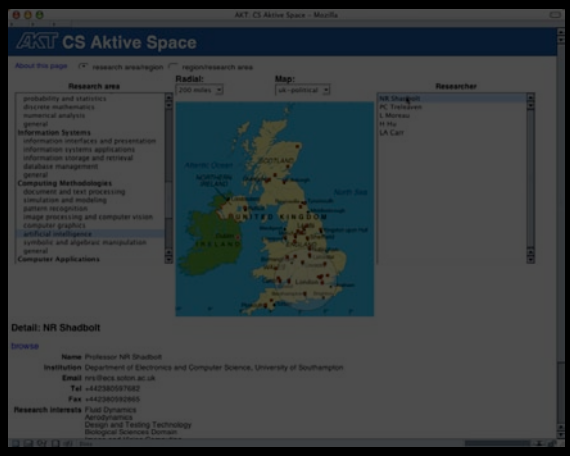

ReSIST - Start Again

- A ReSIST Knowledge Base - The *RKB*

- Project Infrastructure support
- Europe (no longer UK-centric), the World
- Up to date
- Extra subject targets (resilience)
- Browser & platform independent
- Engineer for maintenance
  - Empower partners and other contributors
  - Empower other application builders

ReSIST - and deliver

- D10 - 2007-01-01T00:00:00A
- In fact it is just a URI to a service:
  - http://resist.ecs.soton.ac.uk/sparql/
- Or the raw content can be browsed
  - http://resist.ecs.soton.ac.uk/browse/
- But there is a brand new faceted browser
  - http://resist.ecs.soton.ac.uk/explorer/
- The RKB is embedded in the infrastructure
- The prototype is already being used

Structure

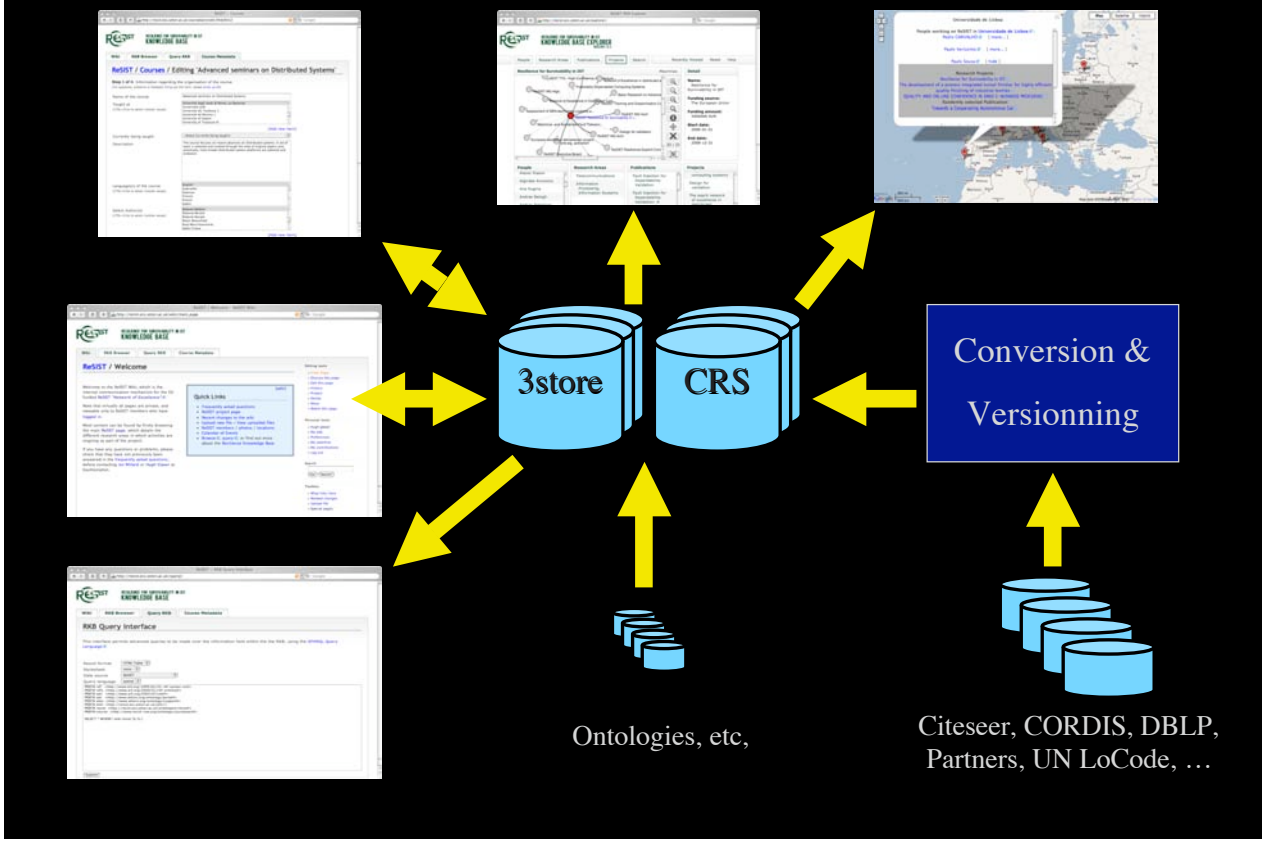

Sources

- Publications
  - Partners
  - Citeseer
  - DBLP
  - ACM
  - DSN & FTCS Series
- Documents
  - RISKS Digest
- Projects
  - CORDIS
  - NSF
- People
  - Partners
- Support
  - UN LoCode

Ontologies etc.

- AKT Ontology
  - Scientific Research Activity
  - Dates
  - Location
  - …
- ALRL Paper
- Courseware (extension of LOM)
- RISKS Codes
- ACM Classification

Main Browser - RKB Explorer

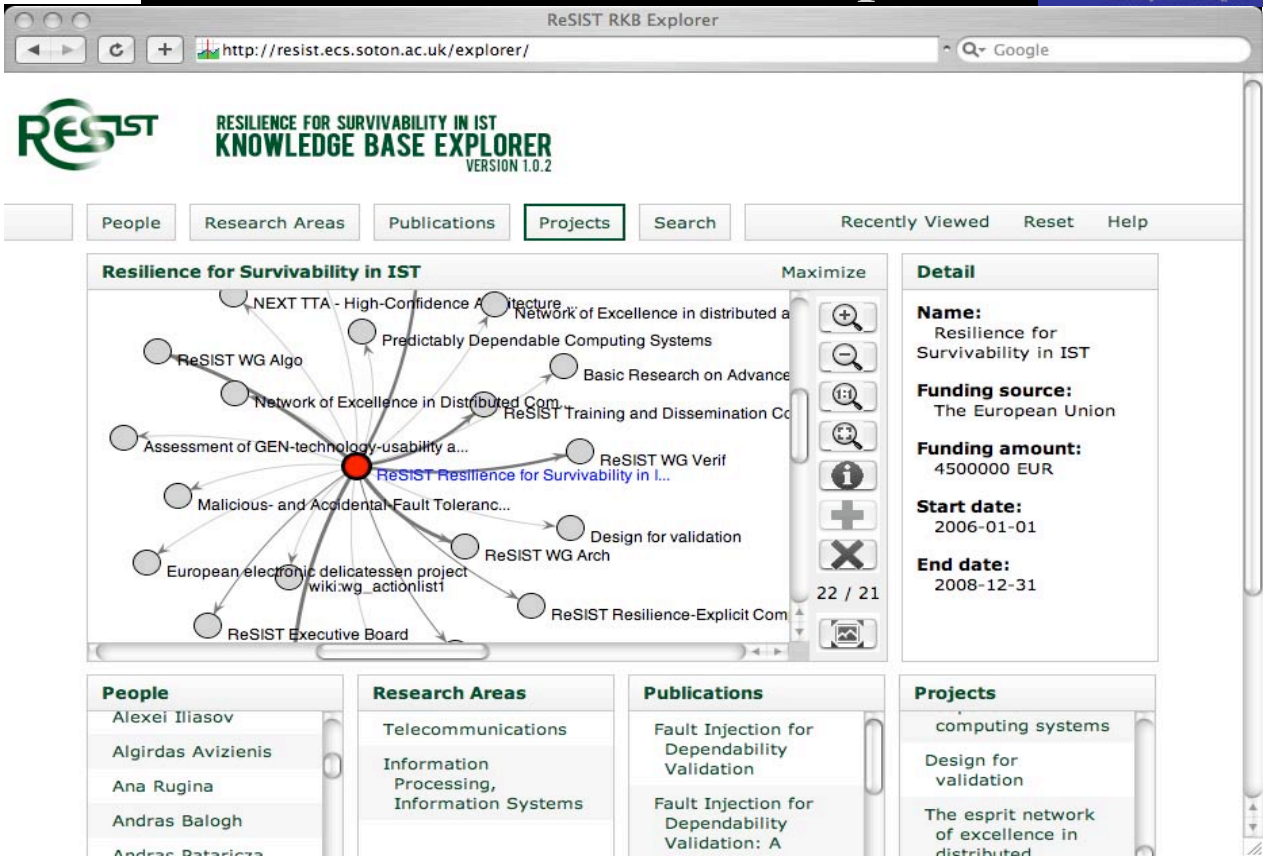

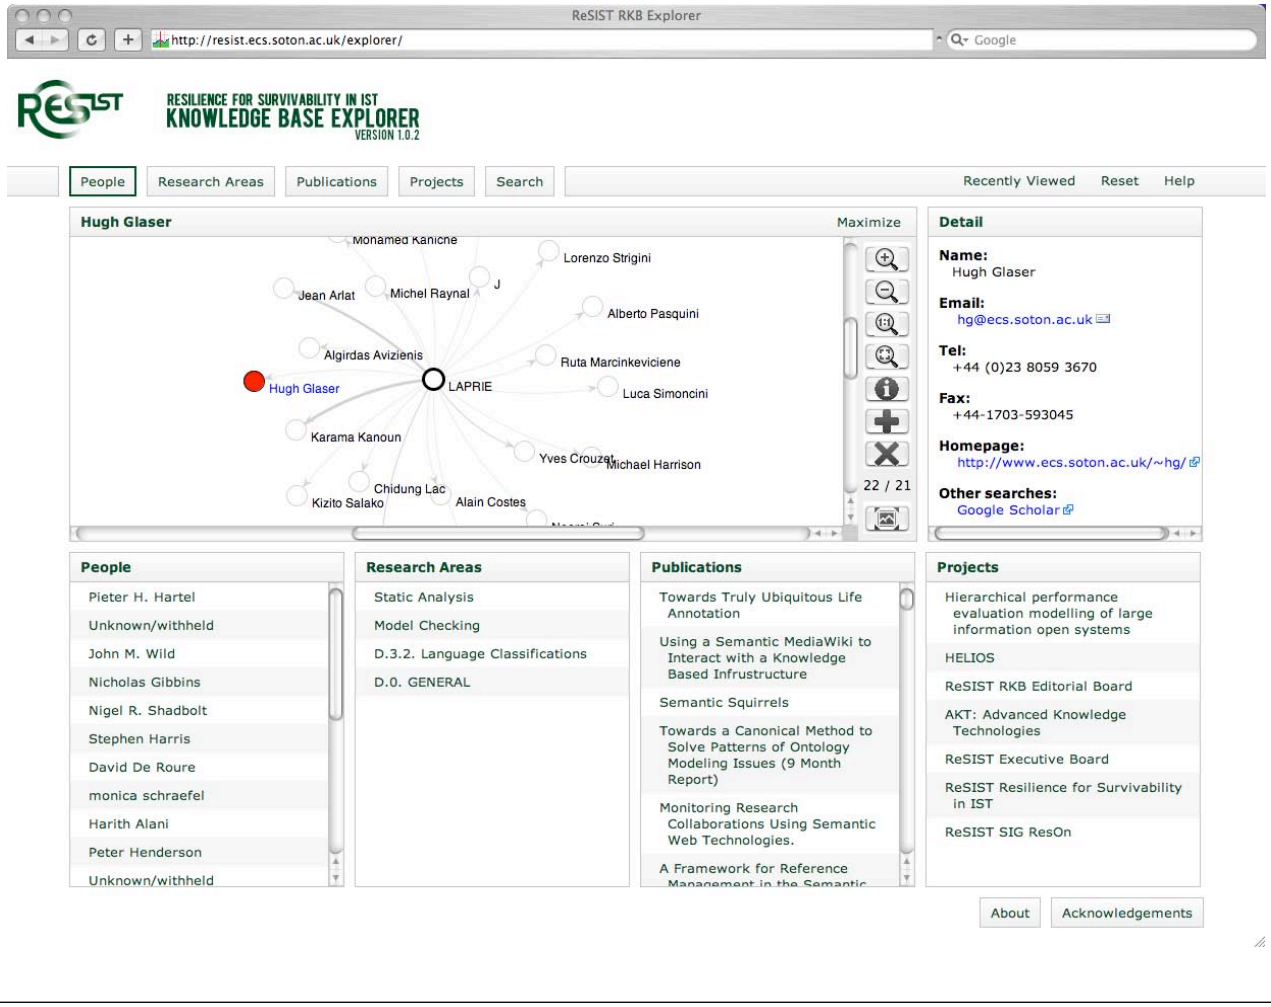
ReSIST RKB Explorer
http://resist.ecs.soton.ac.uk/explorer/
Google
RESILIENCE FOR SURVIVABILITY IN IST
KNOWLEDGE BASE EXPLORER
VERSION 1.0.2
People
Research Areas
Publications
Projects
Search
Recently Viewed
Reset
Help
Hugh Glaser
Maximize
Mohamed Kaniche
Lorenzo Strigini
Jean Arlat
Michel Raynal
J
Alberto Pasquini
Algirdas Avizienis
Ruta Marcinkeviciene
Hugh Glaser
LAPRIE
Luca Simoncini
Karama Kanoun
Yves Crouzet
Michael Harrison
Chidung Lac
Kizito Salako
Alain Costes
22 / 21
Detail
Name:
Hugh Glaser
Email:
hg@ecs.soton.ac.uk
Tel:
+44 (0)23 8059 3670
Fax:
+44-1703-593045
Homepage:
http://www.ecs.soton.ac.uk/~hg/
Other searches:
Google Scholar
People
Pieter H. Hartel
Unknown/withheld
John M. Wild
Nicholas Gibbins
Nigel R. Shadbolt
Stephen Harris
David De Roure
monica schraefel
Harith Alani
Peter Henderson
Unknown/withheld
Research Areas
Static Analysis
Model Checking
D.3.2. Language Classifications
D.0. GENERAL
Publications
Towards Truly Ubiquitous Life Annotation
Using a Semantic MediaWiki to Interact with a Knowledge Based Infrustructure
Semantic Squirrels
Towards a Canonical Method to Solve Patterns of Ontology Modeling Issues (9 Month Report)
Monitoring Research Collaborations Using Semantic Web Technologies.
A Framework for Reference
Projects
Hierarchical performance evaluation modelling of large information open systems
HELIOS
ReSIST RKB Editorial Board
AKT: Advanced Knowledge Technologies
ReSIST Executive Board
ReSIST Resilience for Survivability in IST
ReSIST SIG ResOn
About
Acknowledgements

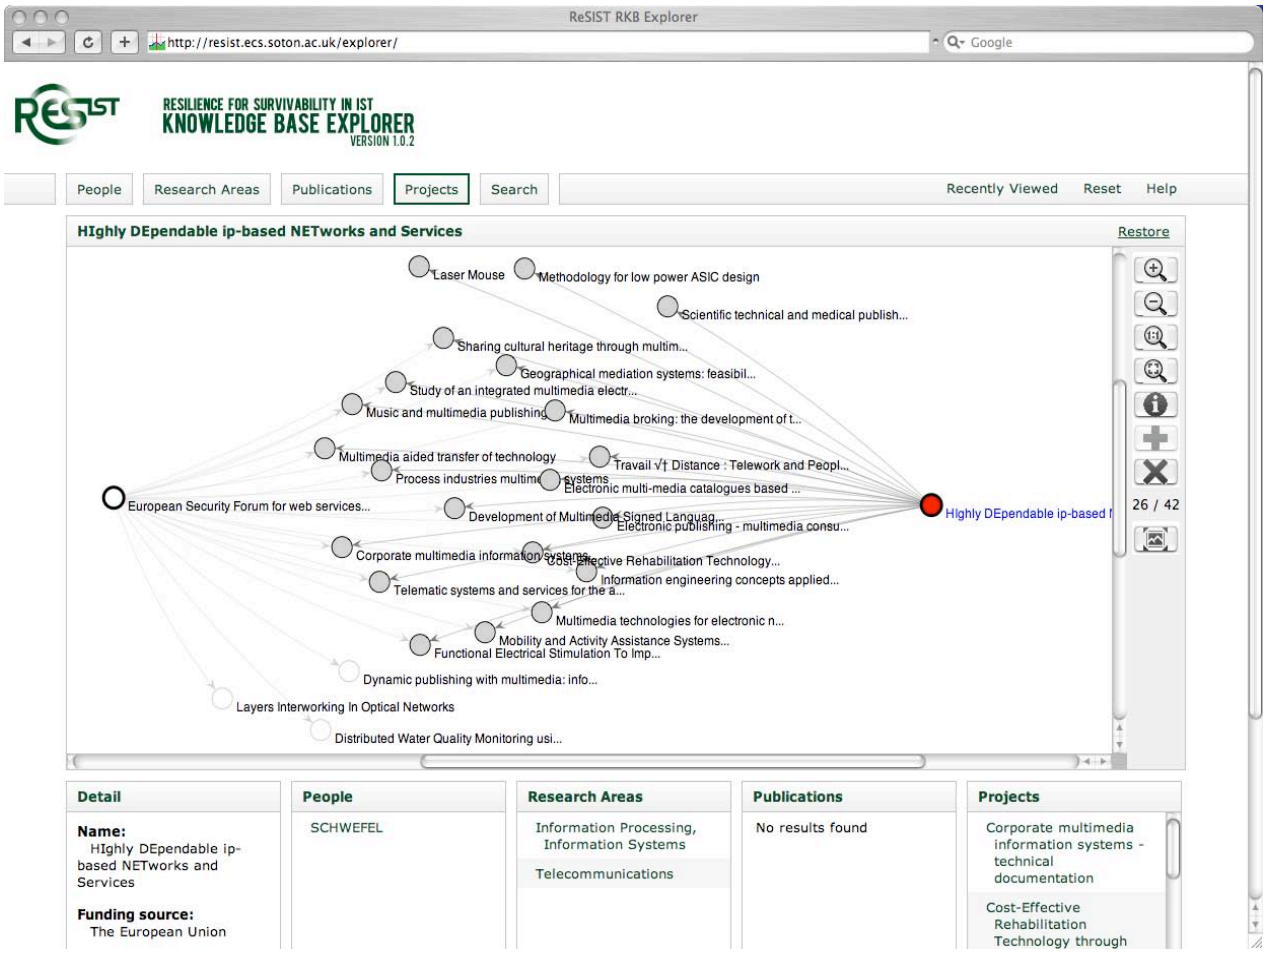
ReSIST RKB Explorer
http://resist.ecs.soton.ac.uk/explorer/
Google
RESILIENCE FOR SURVIVABILITY IN IST
KNOWLEDGE BASE EXPLORER
VERSION 1.0.2
People
Research Areas
Publications
Projects
Search
Recently Viewed
Reset
Help
HIghly DEpendable ip-based NETworks and Services
Restore
Laser Mouse
Methodology for low power ASIC design
Scientific technical and medical publish...
Sharing cultural heritage through multim...
Geographical mediation systems: feasibil...
Study of an integrated multimedia electr...
Music and multimedia publishing
Multimedia broking: the development of t...
Multimedia aided transfer of technology
Travail √† Distance : Telework and Peopl...
Process industries multim
systems
Electronic multi-media catalogues based ...
European Security Forum for web services...
Development of Multimedia Signed Languag...
Electronic publishing - multimedia consu...
HIghly DEpendable ip-based
26 / 42
Corporate multimedia information systems
Cost-Effective Rehabilitation Technology...
Information engineering concepts applied...
Telematic systems and services for the a...
Multimedia technologies for electronic n...
Mobility and Activity Assistance Systems...
Functional Electrical Stimulation To Imp...
Dynamic publishing with multimedia: info...
Layers Interworking In Optical Networks
Distributed Water Quality Monitoring usi...
Detail
Name:
HIghly DEpendable ip-based NETworks and Services
Funding source:
The European Union
People
SCHWEFEL
Research Areas
Information Processing, Information Systems
Telecommunications
Publications
No results found
Projects
Corporate multimedia information systems - technical documentation
Cost-Effective Rehabilitation Technology through

Searching

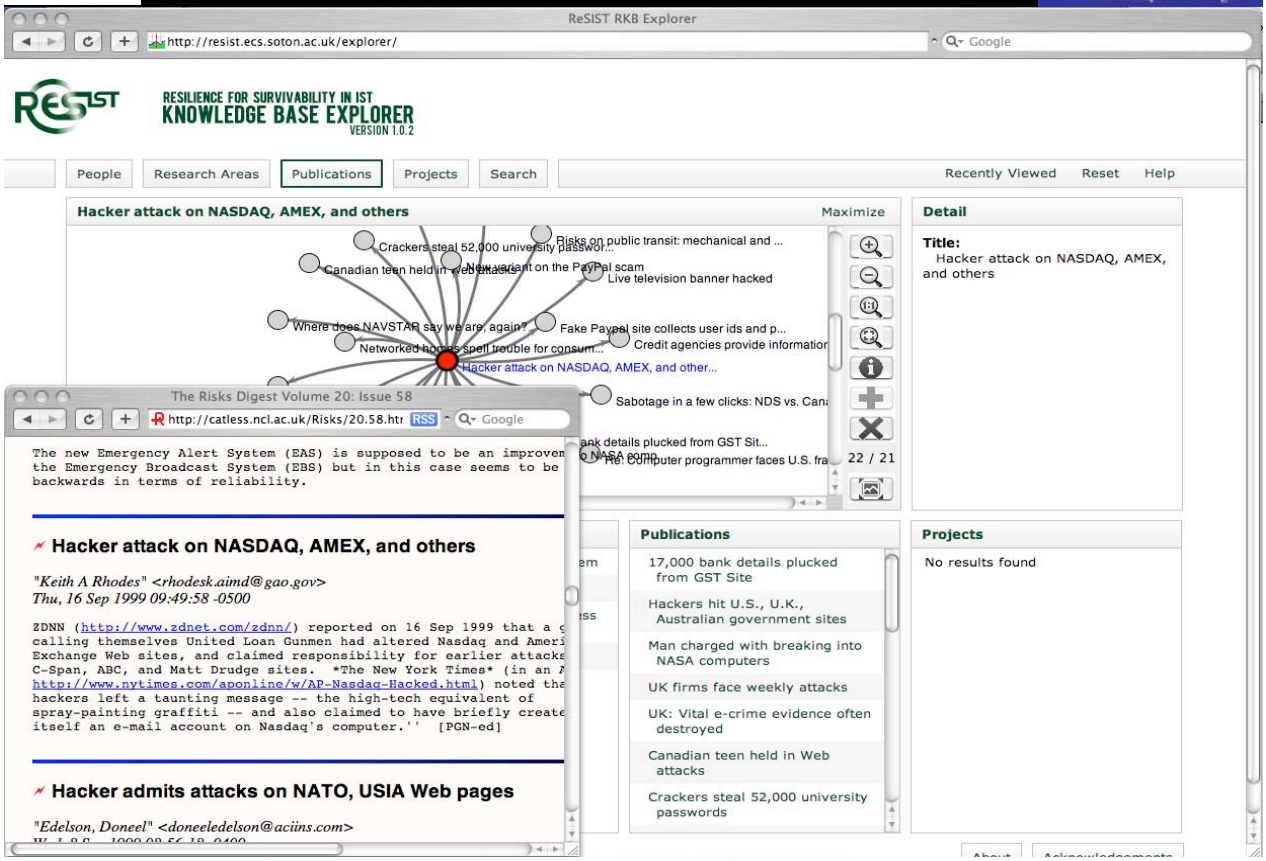

Location

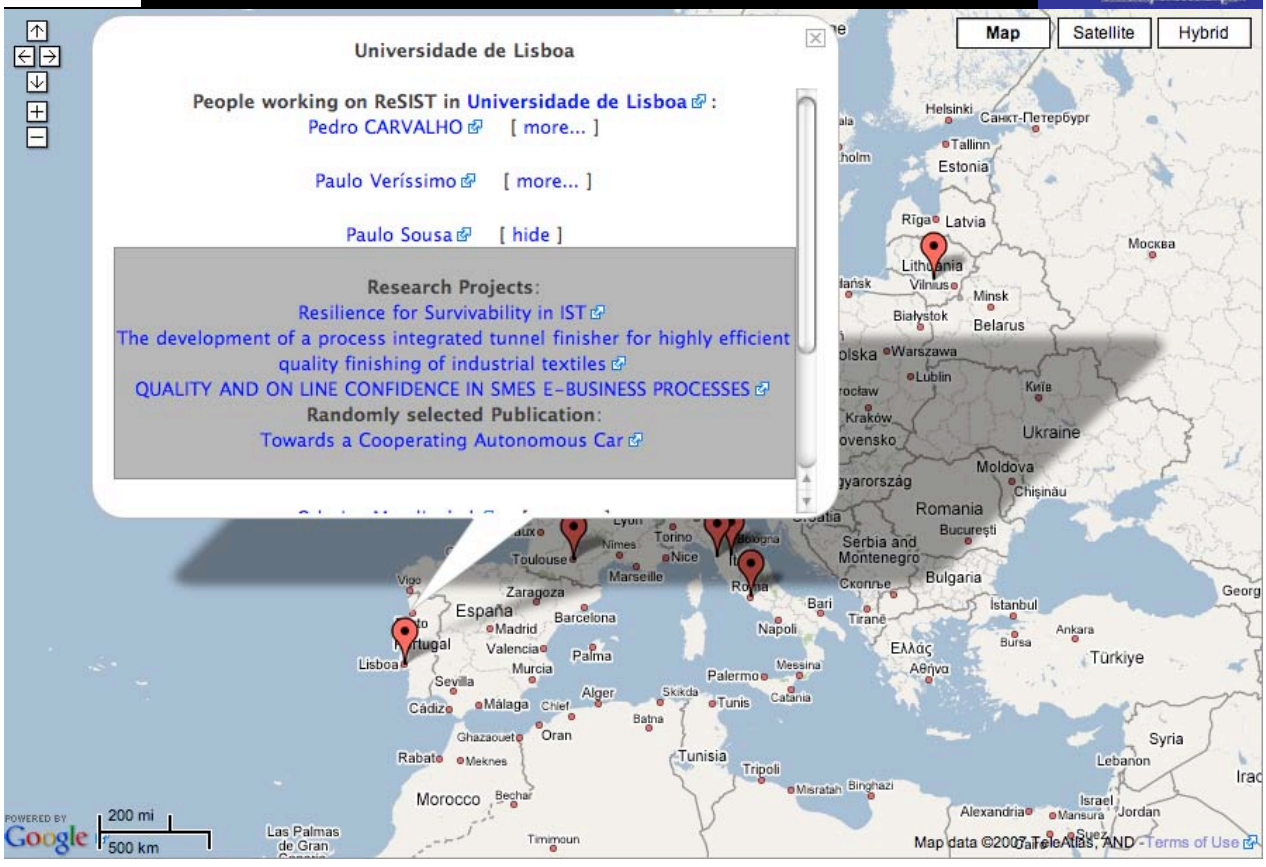

Course Metadata

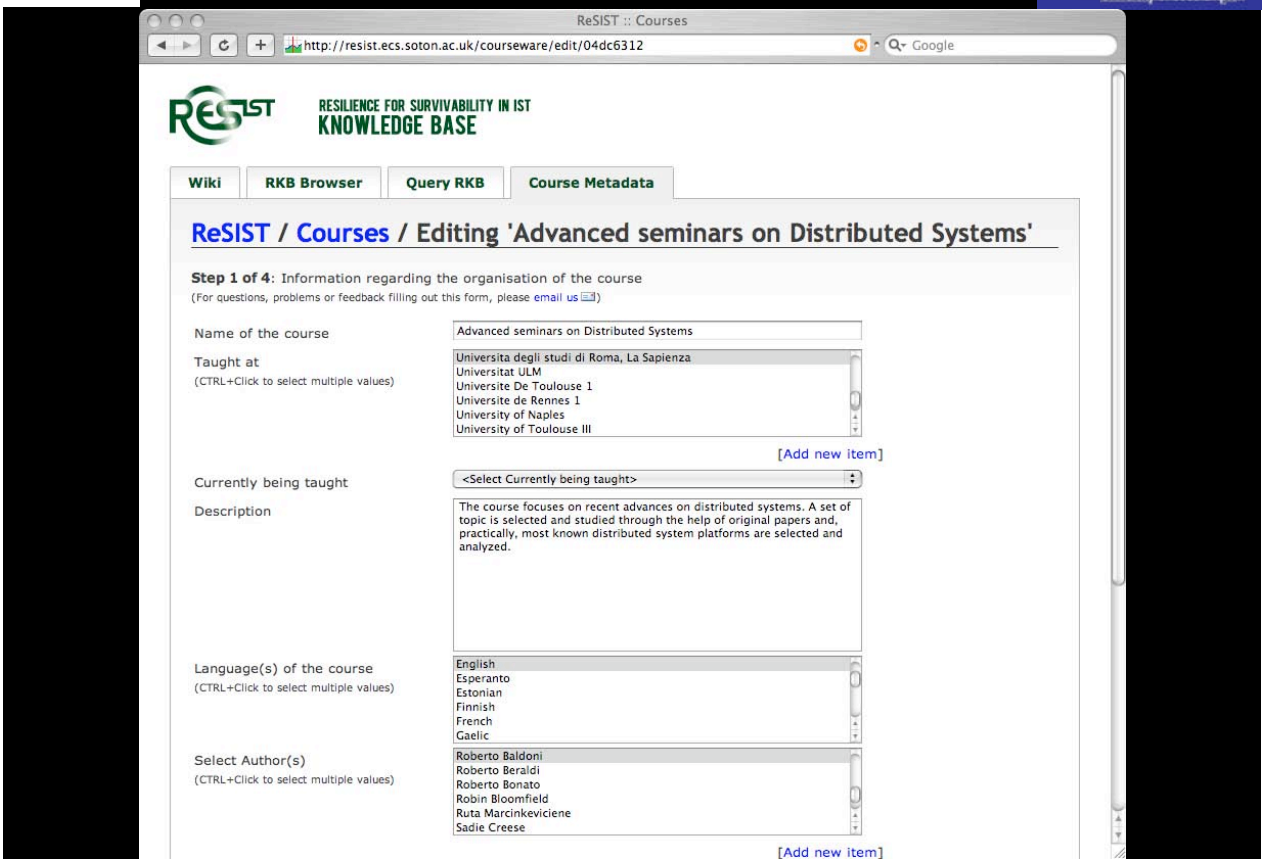

Course Locations

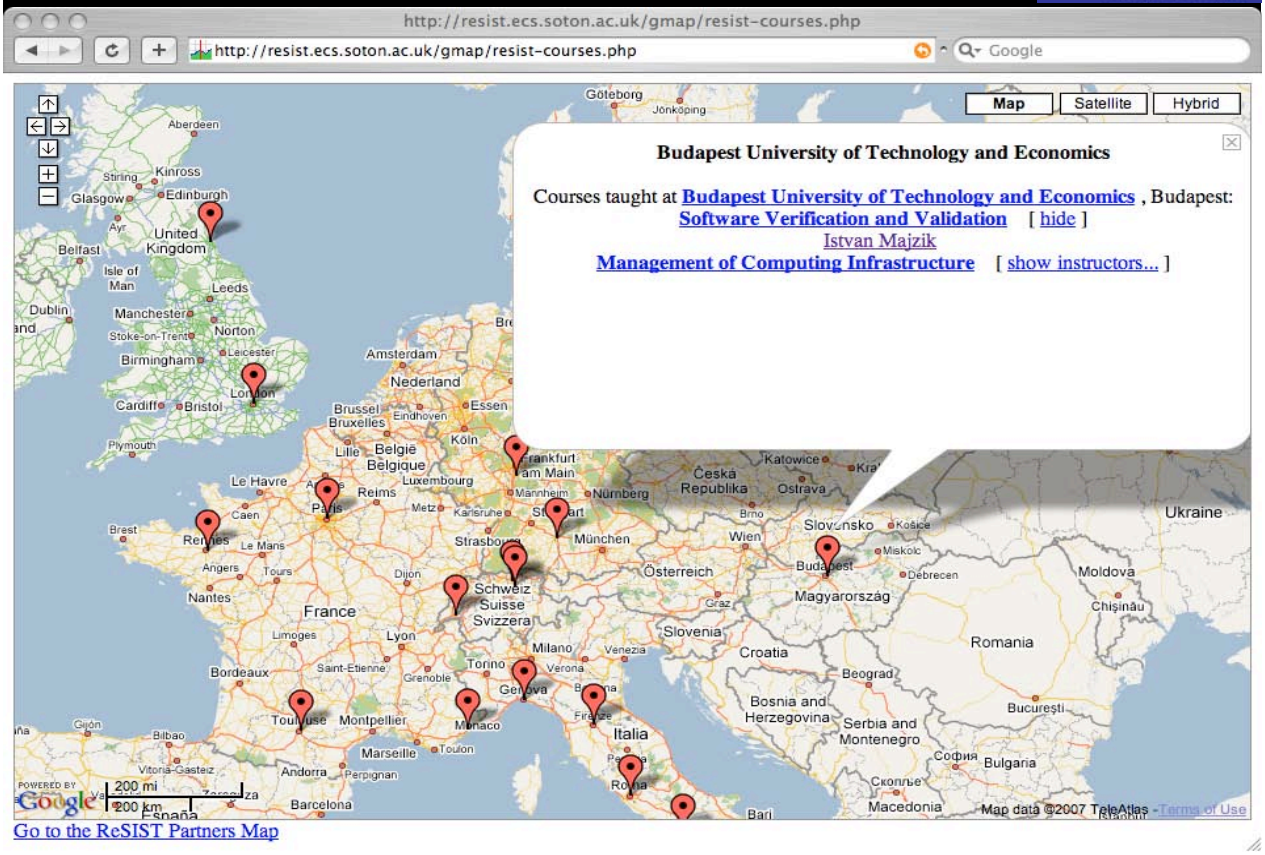

Go to the ReSIST Partners Map

Semantic Wiki

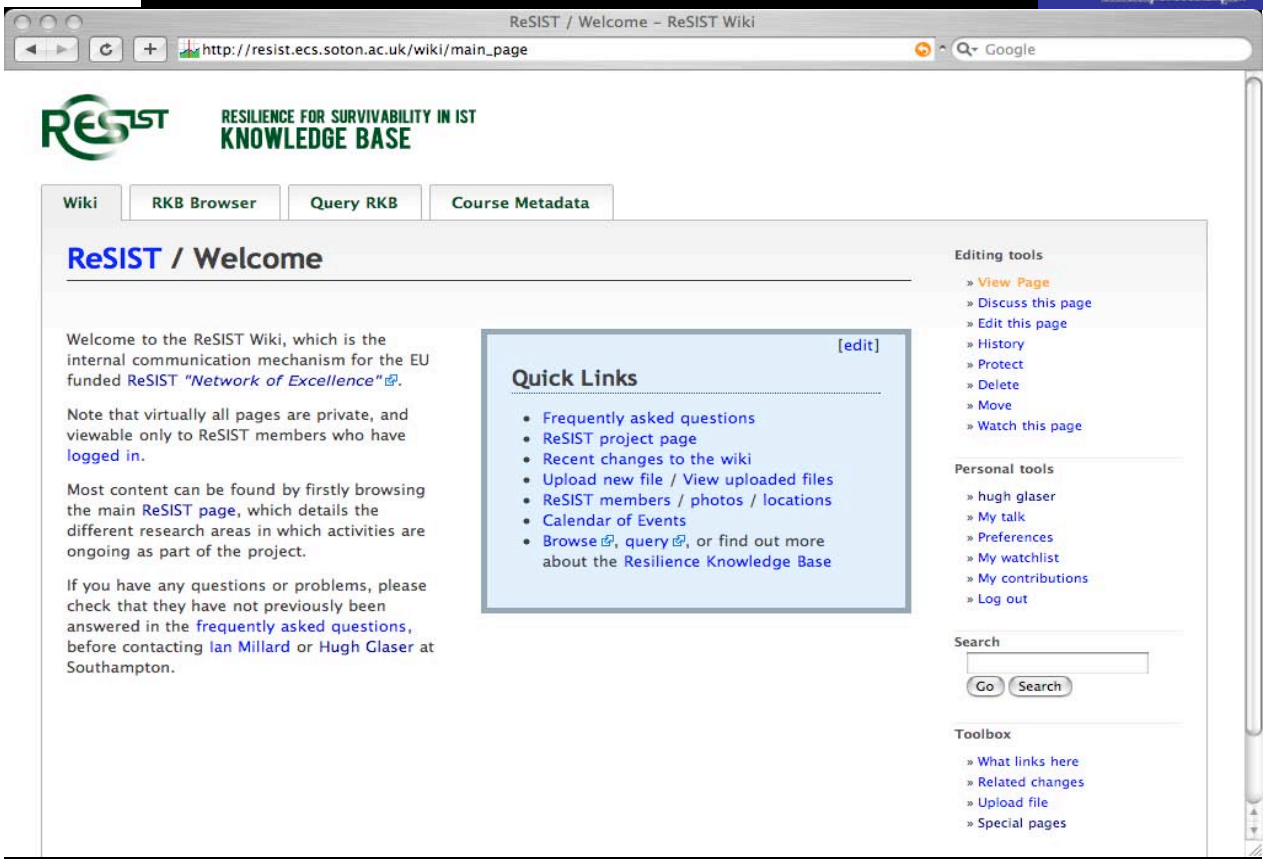

Editing

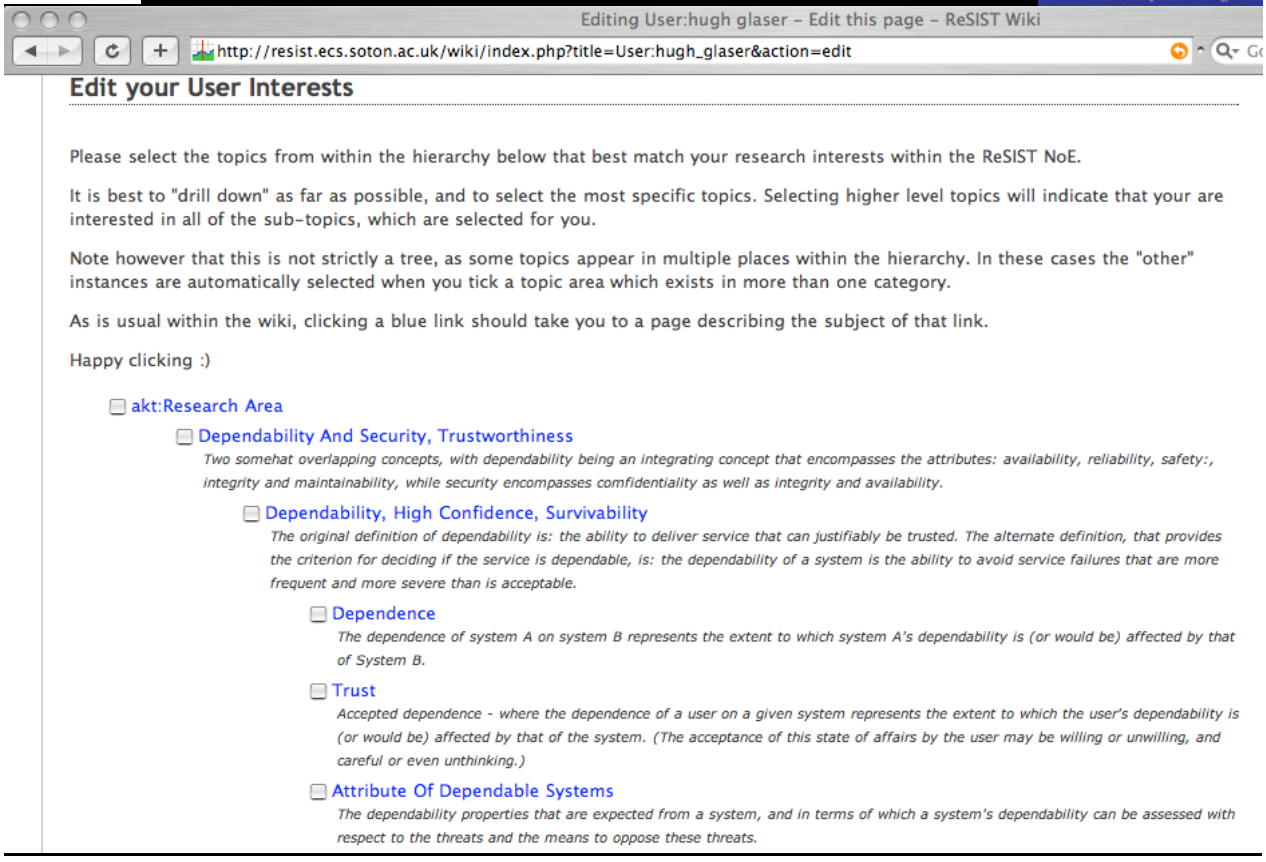

Classifying

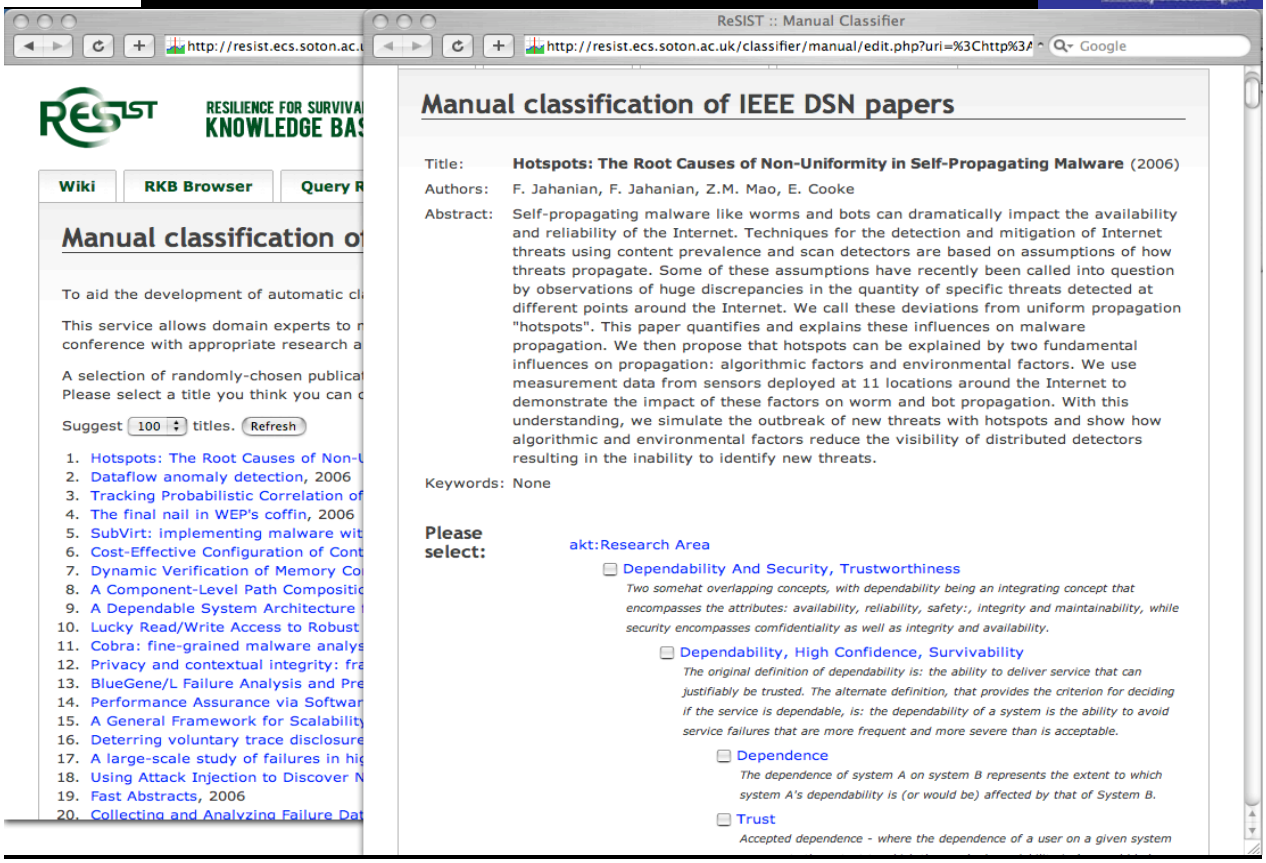

Browsing Raw Data

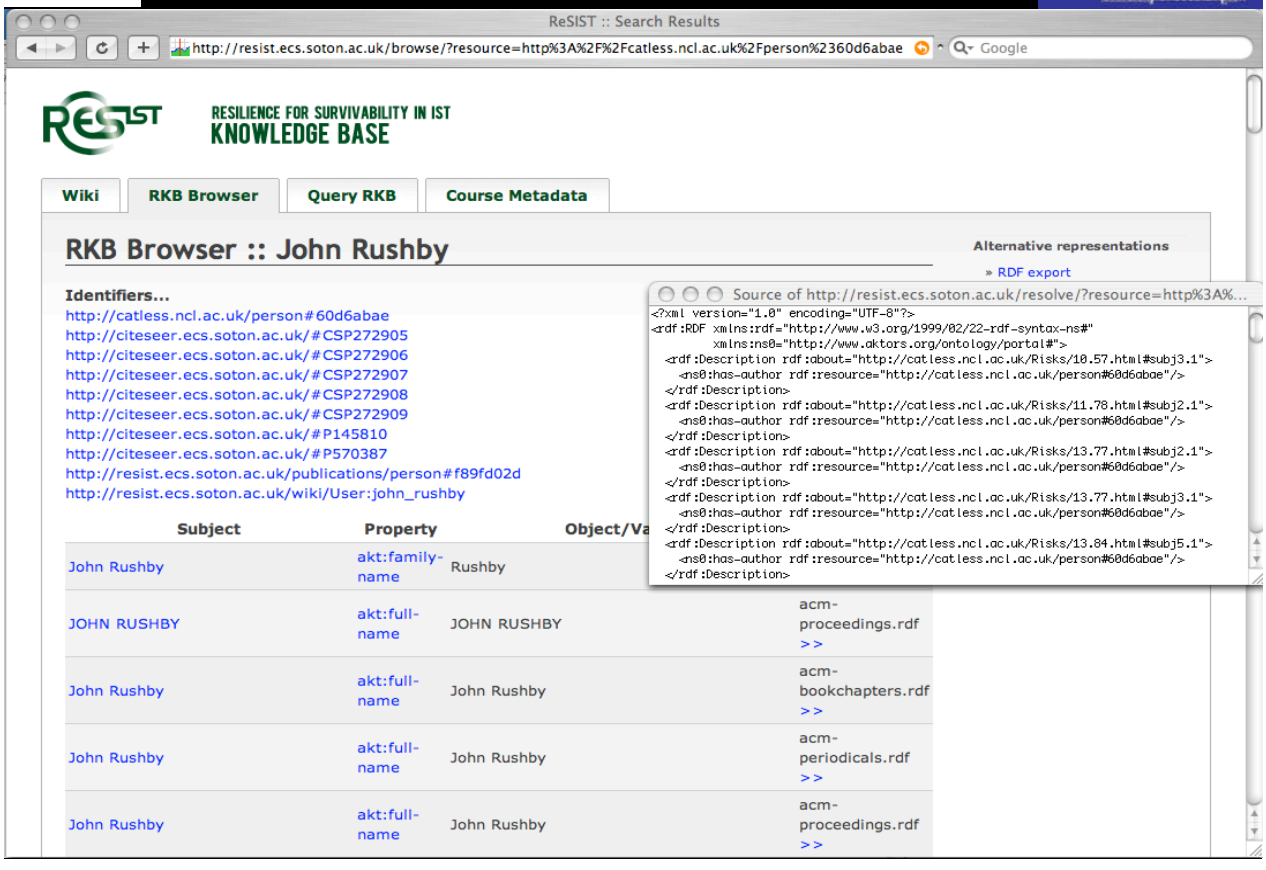

Open Access

ReSIST :: RKB Query Interface

http://resist.ecs.soton.ac.uk/sparql/

RESILIENCE FOR SURVIVABILITY IN IST
KNOWLEDGE BASE

Wiki | RKB Browser | Query RKB | Course Metadata

RKB Query interface

This interface permits advanced queries to be made over the information held within the the RKB, using the SPARQL Query Language.

Result format: HTML Table
Stylesheet: none
Data source: ReSIST
Query language: sparql

```
PREFIX rdf:  <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX owl:  <http://www.w3.org/2002/07/owl#>
PREFIX akt:  <http://www.aktors.org/ontology/portal#>
PREFIX akts: <http://www.aktors.org/ontology/support#>
PREFIX wiki: <http://resist.ecs.soton.ac.uk/wiki/>
PREFIX resist: <http://resist.ecs.soton.ac.uk/ontologies/resist#>
PREFIX course: <http://www.resist-noe.org/ontology/courseware#>

SELECT * WHERE { wiki:resist ?p ?o }
```

Submit

At the Centre

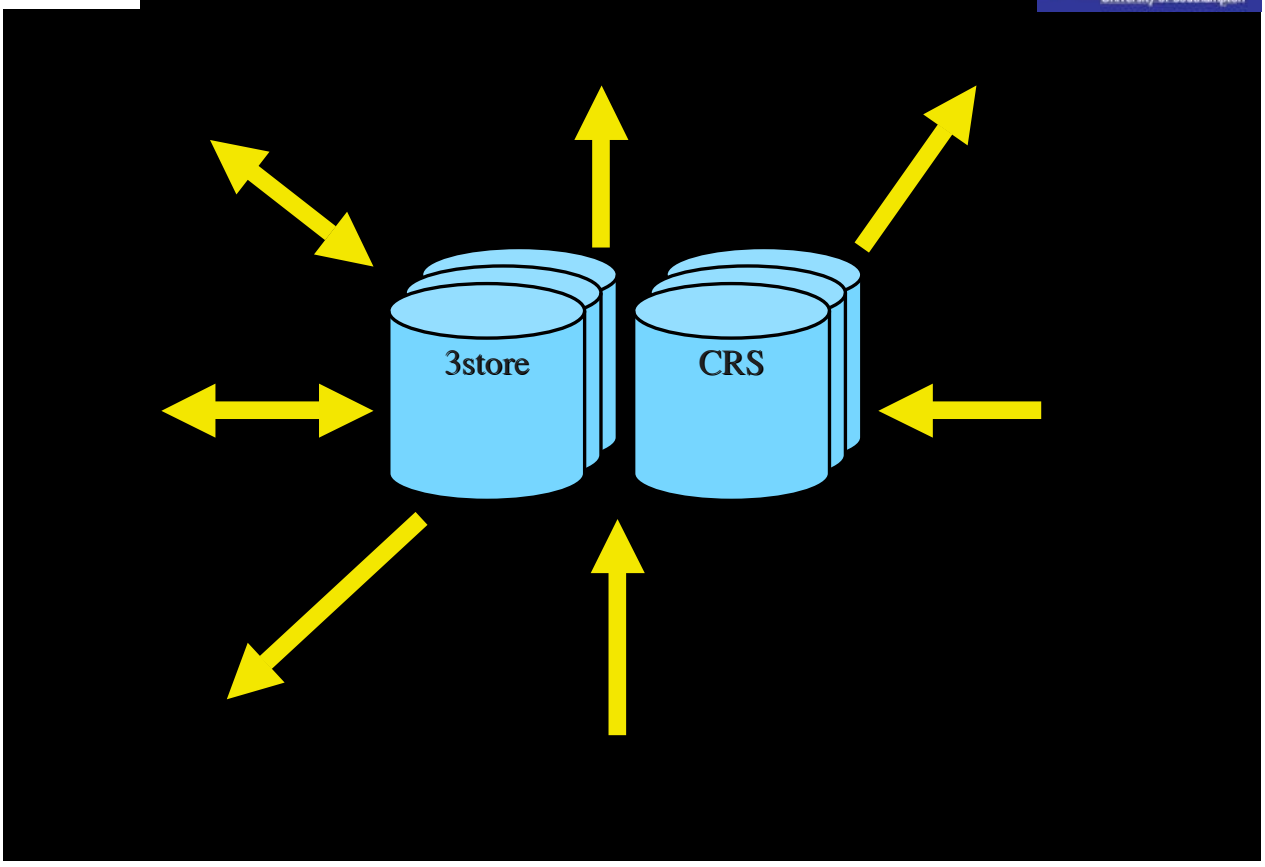

So what *is* RDF...?

- Resource Description Framework
- W3C recommendation
  - From Semantic Web research efforts
- Modelling language
  - Represents facts about resources
- Can model any abstract domain
  - Things do not have to be accessible *on* the web
  - But can be described *in it*

RDF: Basic components

- RDF graphs are formed by *triples*

Important Components 3store and CRS

- 3store
  - Open source semantic store
  - Scalable
    - ReSIST - 50 million facts
      - (cf Wikipedia metadata)
- CRS - Consistent Reference Service
  - Bridges between disparate sources

Openness

- Almost nothing shown was private
- Except
  - Wiki project discussion pages
    - But semantic relations go to RKB
  - Data entry
    - Controlled
    - Not moderated

Future for ReSIST & the RKB

- Improve on the Prototype
  - Sources
  - CRS
  - UI
- Resilient-Explicit Computing
  - Model expert knowledge
  - Model processes, components, mechanisms
- Support Engineer/Scientist
  - Move effectively between
    - System design
    - Knowledge Base
    - People
  - To choose cost, characteristics, etc
- Support Run-Time Deployment
  - Dynamic Reconfiguration

Future Resources

- Original proposal
  - Now primarily maintenance
- Victim of success?
  - Important infrastructure
  - Serious resources to be maintained
  - People want to provide data (costs)

Response

- ReSIST
  - Has increased future RKB resources
- Other Funding and Additionality
  - Lithuania & Saarbrücken
  - JISC
- Longer term
  - Self-funding - SIGs, Clubs
  - Infrastructure - EU, EPSRC, NSF
- Engineer for maintenance and Openess
- Open
  - Knowledge Sources
  - Knowledge Publishing

Some Review Highlights

- One year of work - one RF funded
- ReSIST has done what it said it would do
  - And more
    - In particular, 1M -> 40M
    - Sophisticated UI
- Real tool for the network, from Day One
- Excellent Partner co-operation
  - Data
  - Evaluation
  - Ontology work
- Much Value in Expert Involvement

Modelling of failures: From chains to coincidences

Erik Hollnagel
Professor, Industrial Safety Chair
École nationale supérieure des Mines de Paris, Pôle Cindyniques
Sophia Antipolis, France
E-mail: erik.hollnagel@cindy.ensmp.fr

The future is uncertain

Risk assessment requires an adequate representation – or model – of the possible future events.

The representation must be powerful enough to capture the functional complexity of the system being analysed.

Typical representation: Event tree

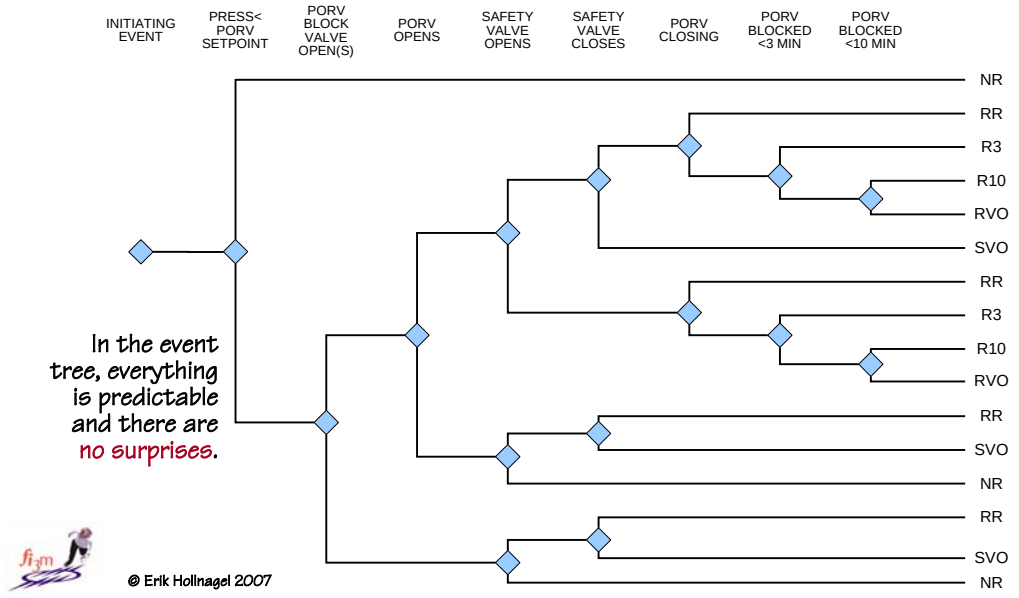

In the event tree, everything is predictable and there are no surprises.

The event tree

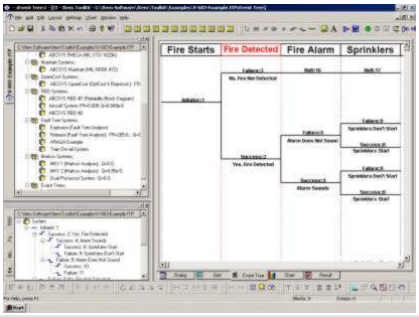

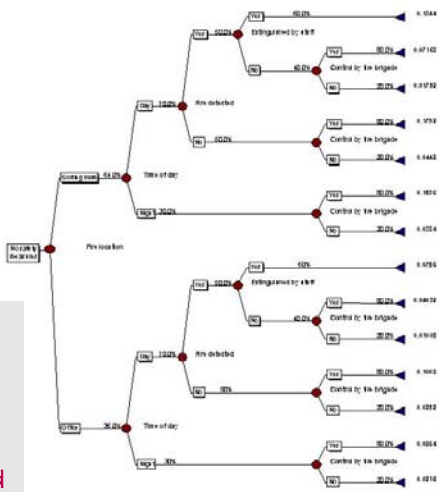

- System is decomposed into elements (components, events)
- Element (failure probability) are described individually
- Element functions are bimodal (true/false, work/fail)
- Order (sequence) is predetermined and fixed
- Linear (non-interacting) combinations
- Limited influence from context/conditions

The fault tree

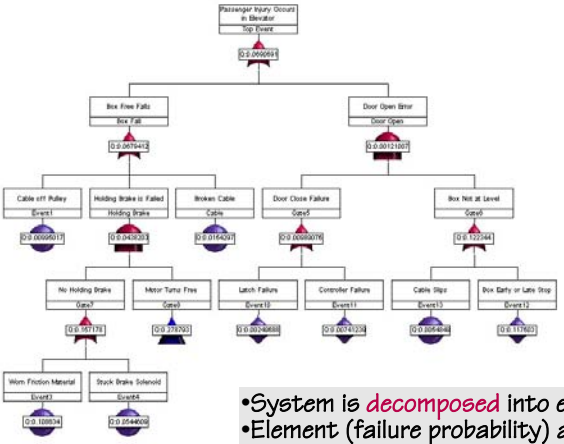

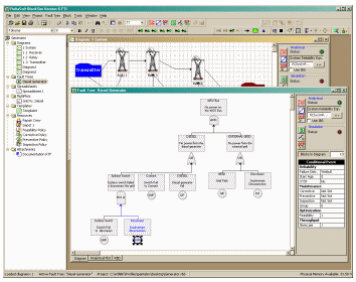

- System is decomposed into elements (components, events)
- Element (failure probability) are described individually
- Element functions are bimodal (true/false, work/fail)
- Order (sequence) is predetermined and fixed
- Linear (non-interacting) combinations
- Limited influence from context/conditions

Nature of technical (formal) systems

Many identical systems

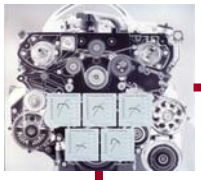

They can be described bottom-up in terms of components and subsystems.

Risks and failures can therefore be analysed relative to individual components and events.

Decomposition works for technical systems, because they have been designed.

Output (effects) are proportional to input (causes) and predictable from knowledge of the components. Technical systems are linear.

Principle of bimodal functioning

In the technological world, things usually function until they fail.
When simple systems, such as a light bulb, fail, they are discarded and replaced by a new (and identical) one.

More intricate systems, such as engines, can be maintained and repaired, as long as it is considered worthwhile.

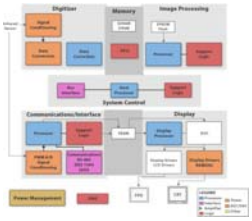

Complex, technological systems work according to the same principle. Failures may, however, be intermittent – especially if complex logic (software) plays a part.
Performance is basically bimodal: either the system works correctly (as designed) or it has failed.

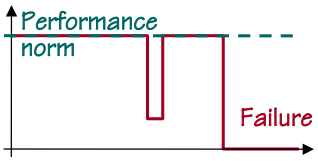

Technological malfunctions

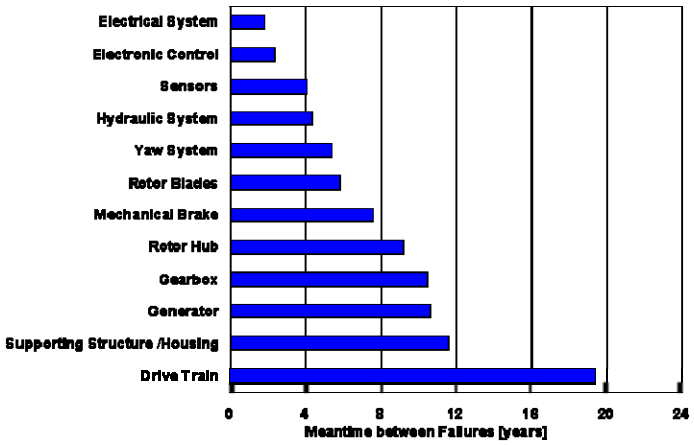

Failure mode
Failure probability
MTBF

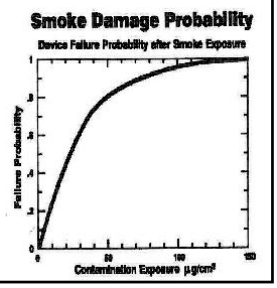

Human malfunctions

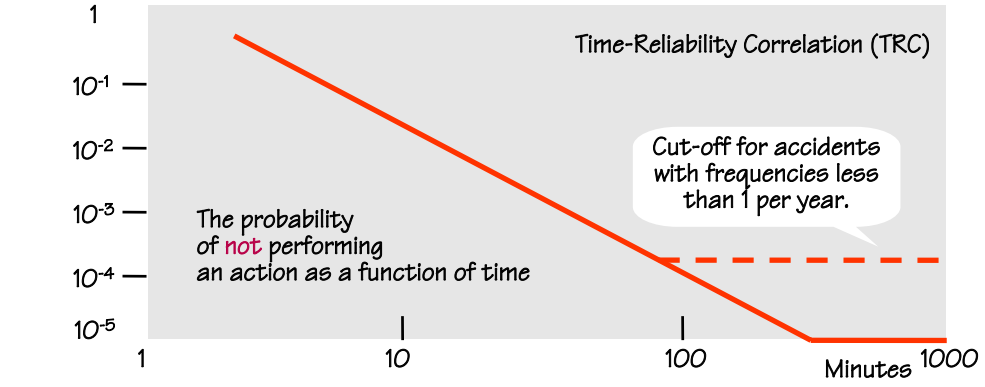

Error of omission (EOO)
Error of commission (EOC)

Failure mode?
Failure probability?
MTBF?

Organizational malfunctions

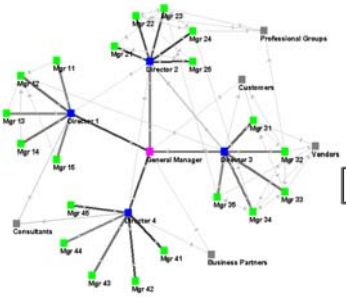

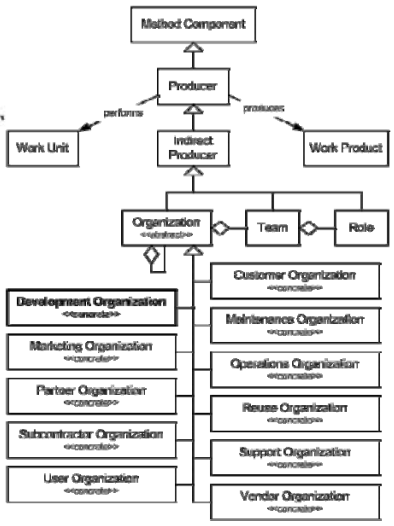

Failure mode?
Failure probability?
MTBF?

Nature of socio-technical systems

All systems unique

Must be described top-down in terms of functions and objectives.

Decomposition does not work for socio-technical systems, because they are emergent.

Risks and failures must therefore be described relative to functional wholes.

Complex relations between input (causes) and output (effects) give rise to unexpected and disproportionate consequences. Socio-technical systems are non-linear.

What is a system?

A system can be defined as "a set of objects together with relationships between the objects and between their attributes" (Hall & Fagen, 1969, p. 81)

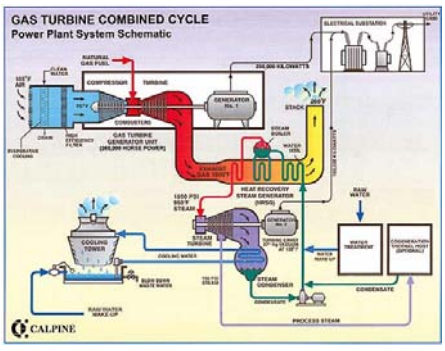

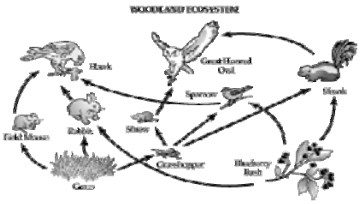

Beer (1964): a manufacturing cell in a garment factory may be considered as a system, as a component of a larger system for garment production, and as containing components, for instance a number of person-cum-scissor units.

There is no 'natural' way of setting the boundary between a system and its environment: it depends on the purpose of the analysis.

Socio-technical systems are not bimodal

Humans and social systems are not bimodal. Normal performance is variable and this – rather than failures and 'errors' – is why accidents happen. Since performance shortfalls are not a simple (additive or proportional) result of the variability, more powerful, non-linear models are needed.

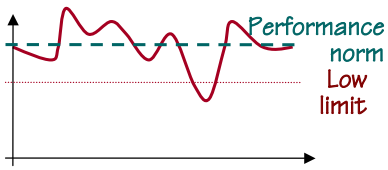

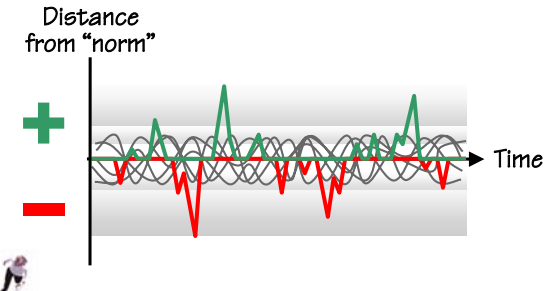

Performance variations can be have positive as well as negative outcomes!

Human factors has tended to look for negative aspects of performance - deviations or "errors"

Traditional view of accidents

The purpose of risk assessment is to identify in a systematic manner how unwanted outcomes can obtain (= severe accidents).

Traditional view:

- Accidents are due to failures or malfunctions of humans or machines. Example: Event Tree
- Risks can be represented by linear combinations of failures or malfunctions. Example: Fault Tree

The chain analogy requires that failures are thought of in a bimodal manner, i.e., something breaks the chain or there is an initial initiating event

Traditional risk assessment is constrained by two assumptions.

Events develop in a pre-defined sequence.

The major source of risk is component malfunctions.

Risk assessment: linear models

Decomposable, simple linear

Sequential accident model → Probability of component failures

Purpose: find the probability that something "breaks", either at the component level or in simple, logical and fixed combinations.
Human failure is treated at the "component" level.

Decomposable, complex linear

Epidemiological accident model → Likelihood of weakened defenses, combinations

Single failures combined with latent conditions, leading to degradation of barriers and defences.

Systemic view of accidents

The purpose of risk assessment is to identify in a systematic manner how unwanted outcomes can obtain (= severe accidents).

Systemic view:

- Accidents are due to unexpected combinations of actions rather than action failures. Example: ETTO.
- Risks can be represented by non-linear combinations of performance variability. Example: FRAM.

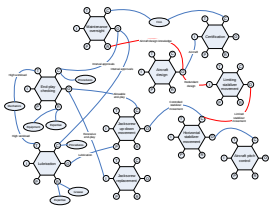

If failures are seen as a result of combinations of normal performance variability rather than of malfunctions, then the chain analogy is no longer adequate.
An alternative approach must be found that emphasises the dynamic nature of how events develop, i.e., coincidences rather than chains.
One possibility is to use resonance rather than failure.

Normal behaviour is variable

Social-technical system failures cannot be modelled as deviations from required or normal performance:
- humans are not designed.
- conditions of work are usually underspecified
- humans are multifunctional, and can do many different things

Accounting for the sources and range of normal performance variability:

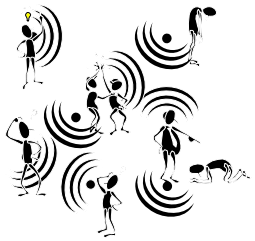

Inherent variability (psychological / physiological phenomena).
Ingenuity and creativity – adaptability (overcoming constraints and underspecification).
Organizationally induced performance variability (meeting demands, stretching resources).
Socially induced variability (meeting expectations, informal work standards).
Contextually induced performance variability (performance conditions).

Risk assessment: non-linear models

Performance variability is natural in socio-technical systems, and a valuable part of normal performance. The many small adjustments enable humans to cope with the complexity and uncertainty of work.

The adjustments allow the system to achieve its functional goals more efficiently by sacrificing details that under normal conditions are unnecessary. Humans are adept at developing working methods that allow them to take shortcuts, thereby often saving valuable time.

Accounting for how performance variability may combine:

Functional resonance (unintended, non-linear outcomes of normal performance adjustments).
Actions based on expectations (of what others have done or will do)
Unanticipated consequences (exact predictions impossible)
Combinations of "unsafe" actions and latent conditions

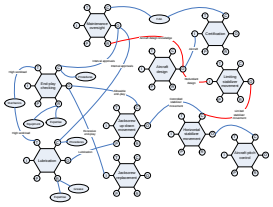

Traffic and randomness

Traffic is a system in which millions of cars every day move so that their driving paths cross each other and critical situations arise due to pure random processes:
cars meet with a speed difference of 100 to more than 200 km/h, separated only by a few meters, with variability of the drivers' attentiveness, the steering, the lateral slope of the road, wind and other factors.

Drivers learn by experience the dimensions of the own car and of other cars, how much space is needed and how much should be allocated to other road users, the maximum speed to approach a curve ahead, etc. If drivers anticipate that these minimum safety margins will be violated, they will shift behavior.

The very basis of traffic accidents consists of random processes, of the fact that we have complicated traffic system with many participants and much kinetic energy involved.
When millions of drivers habitually drive at too small safety margins and make insufficient allowance for (infrequent) deviant behavior or for (infrequent) coincidences, this very normal behavior results in accidents.

Summala (1985)

Looking back only ONE thing happened

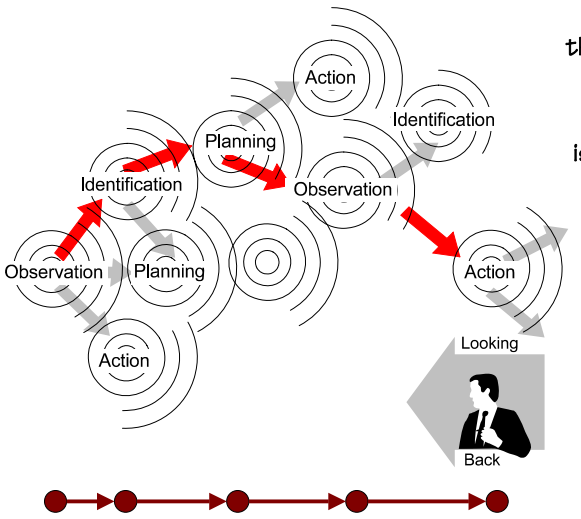

Given the actual context, the events seem to describe an orderly sequence.

The order (chain of events) is, however, an artefact due to the asymmetry of time

Looking ahead ANYTHING can happen

Prediction that is not constrained, is basically a combinatorial effort. The results therefore represent the complexity of the classification system, rather than real performance.

Actions are more often determined by the final cause (telos) than by the efficient cause.
Causal chains are thus of an a posteriori rather than an a priori nature.

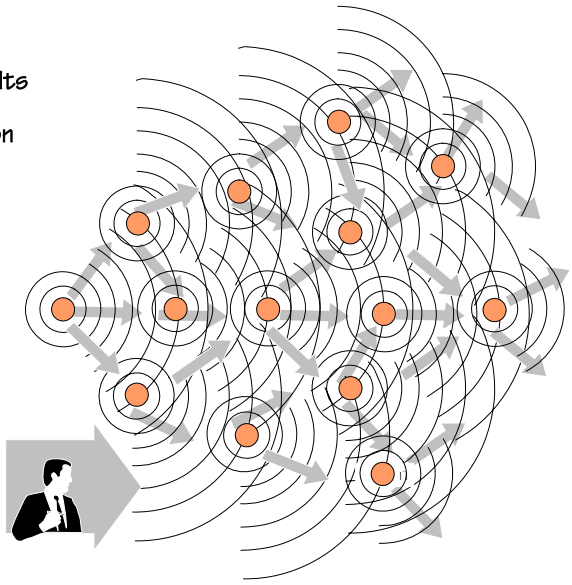

The future as non-linear events

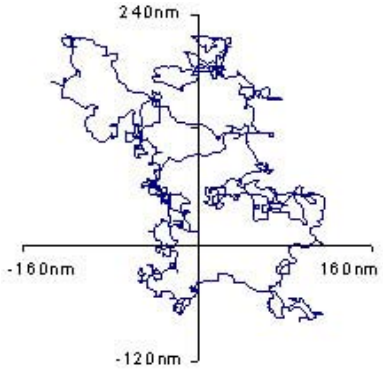

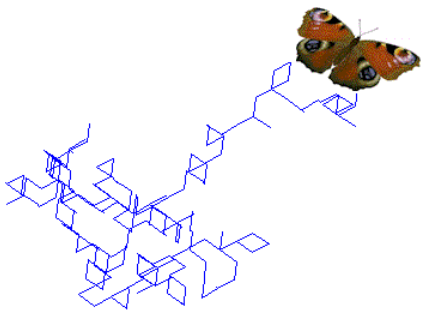

Non-linear events have been likened to Brownian movements or random walks.
Risk assessment requires something that is non-linear (non-trivial) at the same time as it is systematic (predictable)

Resonance

Stochastic resonance

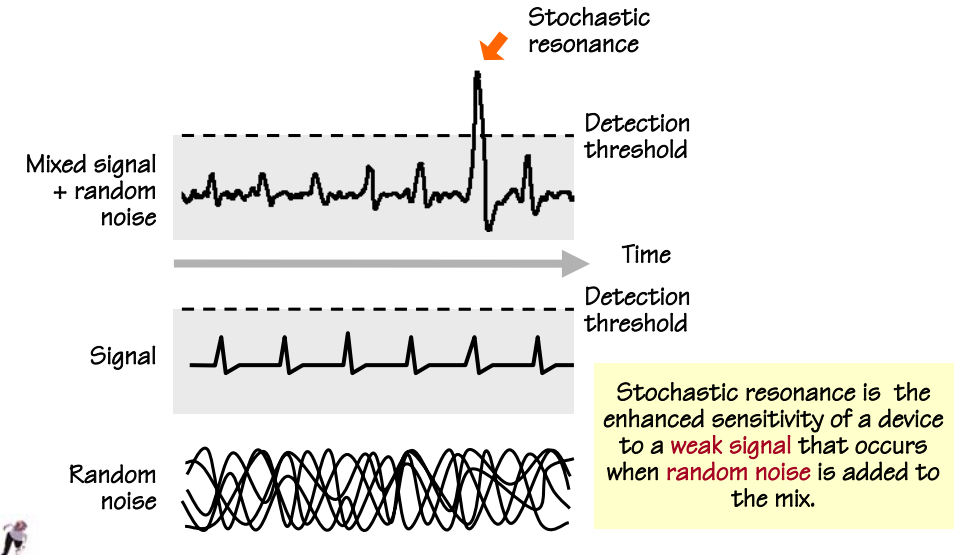

Stochastic resonance is the enhanced sensitivity of a device to a weak signal that occurs when random noise is added to the mix.

Functional resonance

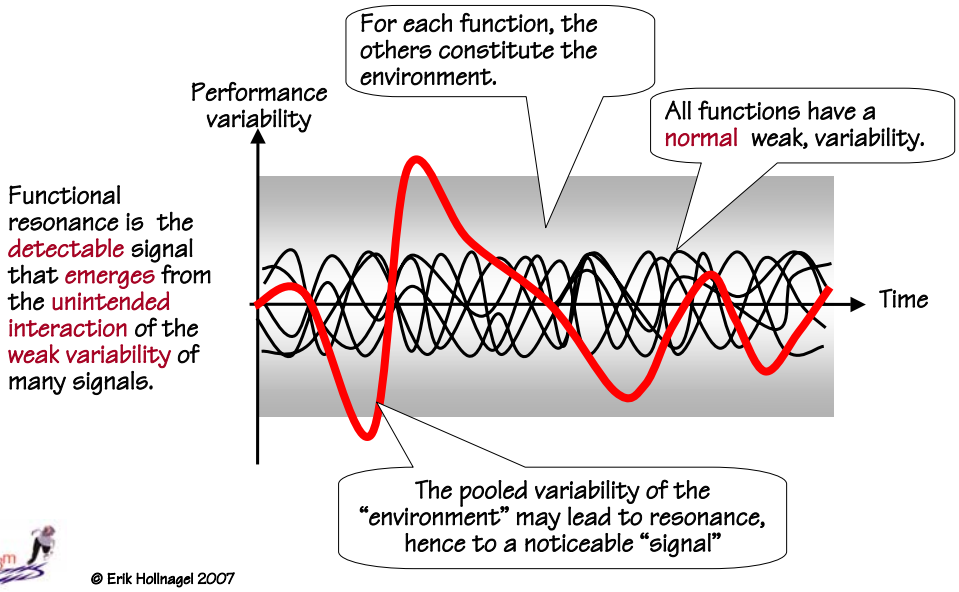

Functional resonance is the detectable signal that emerges from the unintended interaction of the weak variability of many signals.

Functional Resonance Accident Model

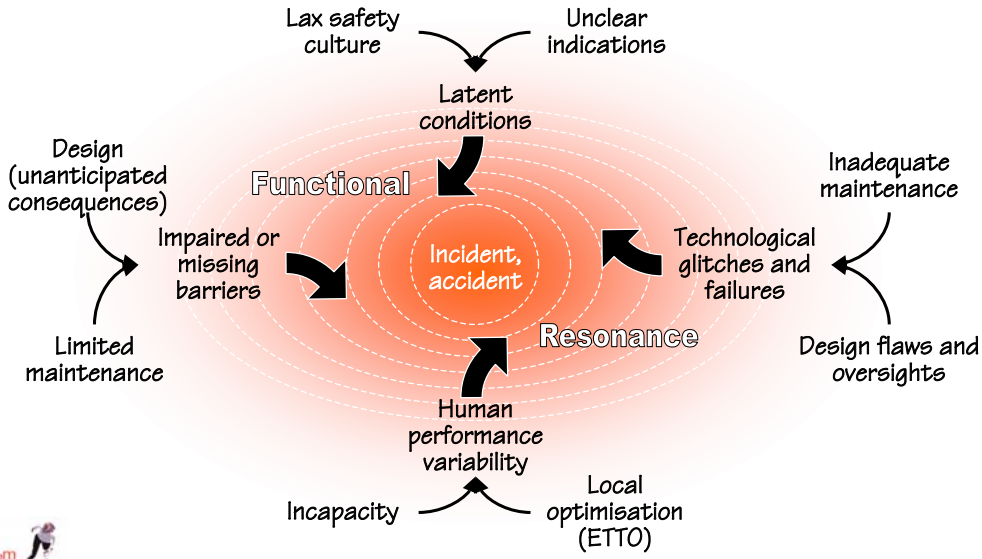

Handling drug prescriptions (HTA)

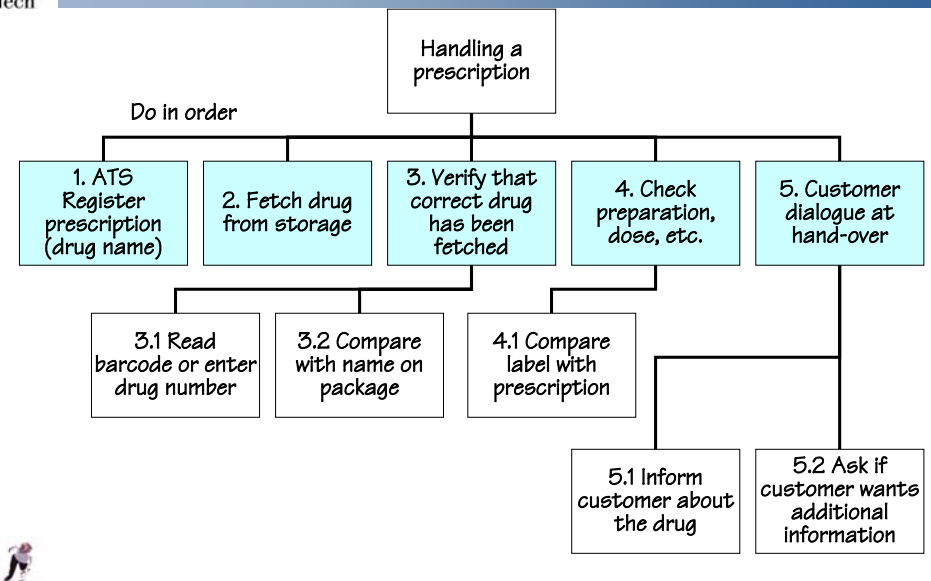

Drug handling – normal procedure

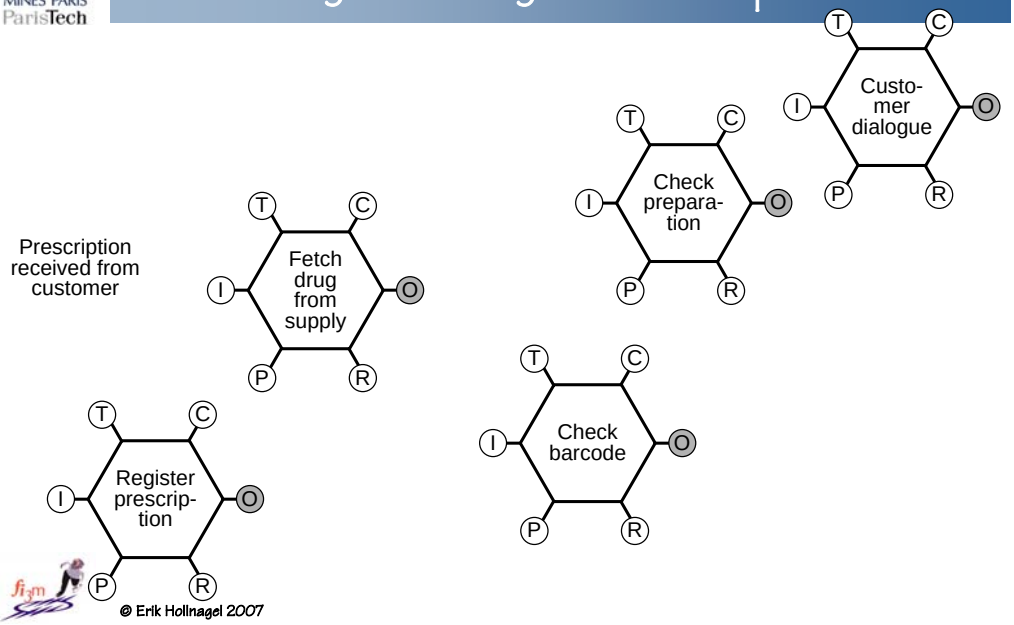

Conclusions

Risk assessment must comprise a model of the system and its behaviour, which is as complex as the system itself.

- Conventional risk assessment is based on linear models (e.g., event tree) and on calculating failure probabilities.
- Socio-technical systems are non-linear. Risk is an emergent rather than a resultant phenomenon.

Risk assessment should address how irregularities can arise from normal performance variability, rather than on how individual functions fail.

- Performance variability reflects the nature of the work environment, including social and organisational factors.
- Performance variability is predictable for identified conditions.

The principle of functional resonance can be used to identify possible combinations of performance variability which may lead to the occurrence of undesirable outcomes.

Three premises of resilience engineering

- **Performance conditions are always underspecified.**
  It is impossible to specify in every detail what should be done and how. Individuals and organisations must therefore always adjust their performance to the current conditions; and because resources and time are finite, such adjustments will inevitably be approximate.
  Performance variability is unavoidable, but it is a source of successes as well as of failures.
- **Many adverse events can be attributed to a breakdown or malfunctioning of components and normal system functions, but many cannot.**
  These are best understood as the result of unexpected combination of normal performance variability. Adverse events therefore represent the converse of the adaptations necessary to cope with the complexity of the real world.
- **Effective safety management cannot be based on hindsight, nor rely on error tabulation and the calculation of failure probabilities.**
  Safety management must be proactive as well as reactive. Resilience Engineering looks for ways to enhance the ability of organisations to create processes that are robust yet flexible, to monitor and revise risk models, and to use resources proactively in the face of disruptions or ongoing production and economic pressures.

Resilience engineering

- Resilience requires an organisation that at all times is:

  Responsive - able to respond effectively when something happens
  Attentive - knows what to look for and regularly updates its knowledge, competence and resources
  Looking ahead - prepared for what might conceivably happen in the future in both the short and the long term.

- The development and application of Resilience Engineering requires

  The ability to measure, monitor, and analyse the resilience of an organisation in its operating environment,
  Tools and methods to improve an organisation's resilience vis-à-vis the environment, and finally
  Techniques to model and predict the short- and long-term effects of changes to operational, organisations, and targets..

- The purpose of safety management is not to reduce risks or the number of adverse events, but to increase on all levels the ability to adjust performance in the face of changes, disturbances, and uncertainty.

Drug handling – normal procedure

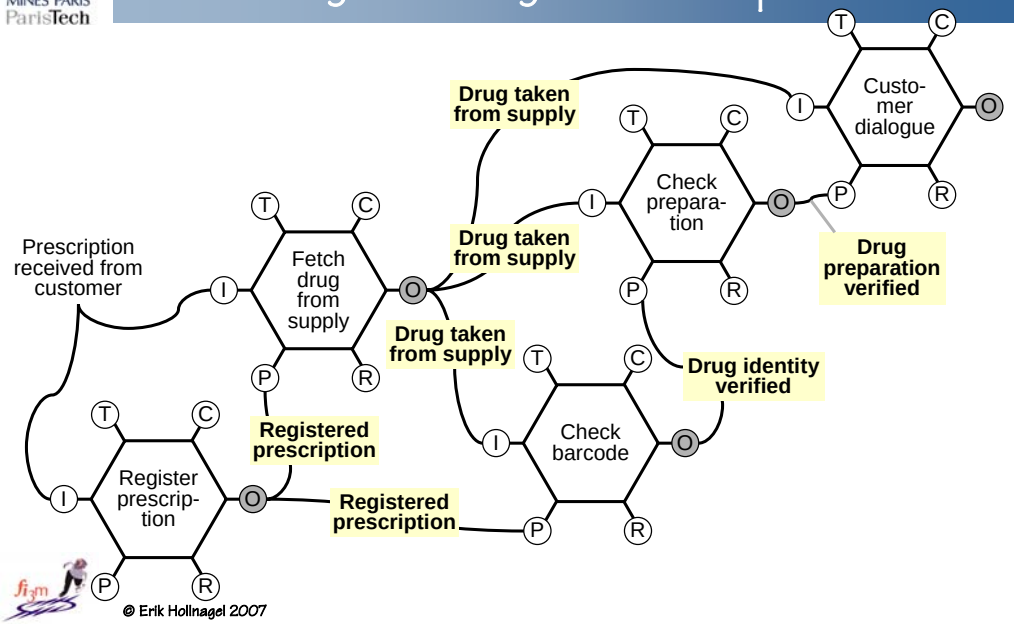

Drug handling - variation

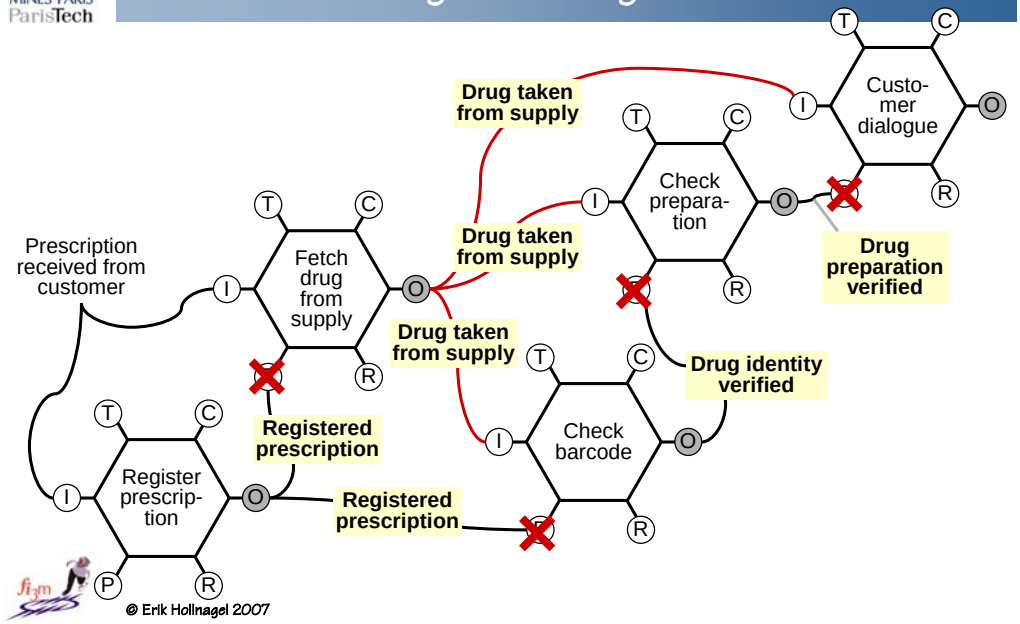

Important announcement

EUROCONTROL

European Organisation for the Safety of Air Navigation

DFS Deutsche Flugsicherung

MINES PARIS ParisTech

Ph.D. Position
"A resilience based approach to evaluate the human contribution to system safety"

The position is part of a new project in a collaboration between Eurocontrol, Deutsche Flugsicherung (DFS), and École des Mines de Paris, Pôle Cindyniques.
The main place of work will be Sophia Antipolis, France

For further information please contact either:
erik.hollnagel@cindy.ensmp.fr
Oliver.Straeter@eurocontrol.int

© Erik Hollnagel 2007

Thank you for your attention

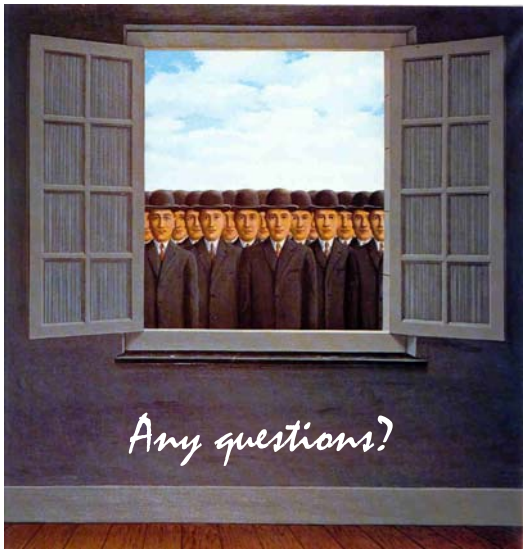

Panel on

Resilience Views from other European Projects

Panel Moderator: **Luca Simoncini**, University of Pisa, Italy - **ReSIST NoE**

Panellists:

Benoît Bruyère, Thales, France - **DESEREC IP**

Aljosa Pasic, Atos Origin, UK - **ESFORS CA**

Domenico Presenza, Engineering Ingegneria Informatica, Italy - **SERENITY IP**

Hans-Peter Schwefel, Aalborg University, Denmark - **HIDENETS STREP**

Information Society Technologies

ReSIST NoE

Resilience for Survivability in IST

SIXTH FRAMEWORK PROGRAMME

DEpendability and Security by Enhanced REConfigurability

DESEREC is an IP of FP6. It deals with highly interconnected Communications and Information Systems (CIS), and the use of them to carry out critical activities. It aims at the development of model-based reconfiguration techniques for large IT systems, thus protecting services against faults and intrusions.

European Security Forum for WEB Services, Software and Systems

ESFORS is a CA of FP6. It aims at bringing together the European stakeholders for security and dependability Information and Communication Technologies (ICTs) to address the security and dependability requirements of emerging software service platforms.

SERENITY System Engineering for Security & Dependability

System Engineering for Security & Dependability

SERENITY is an IP of FP6. It aims to enhance security and dependability in AMI systems, by validated security solutions available to AmI ecosystems and promoting their assurance and evolution. It will provide mechanisms for monitoring security at run-time and dynamically react to threats or breaches of security, and context changes and it will integrate security solutions, requirements definition and solution selection, and monitoring and reaction mechanisms in a common framework.

HIDENETS

HIghly DEpendable ip-based NETworks and Services

HIDENETS is a STREP in FP6. The aim of HIDENETS is to develop and analyze end-to-end resilience solutions for distributed applications and mobility-aware services in ubiquitous communication scenarios. Technical solutions will be developed for applications with critical dependability requirements in the context of selected use-cases of ad-hoc car-to-car communication with infrastructure service support.

Information Society Technologies

ReSIST NoE

Resilience for Survivability in IST

SIXTH FRAMEWORK PROGRAMME

What is "Resilience" ?

Resilience* and **Resilience Engineering*** are defined as:

- in **Networks**: **Resilience** is the ability of the network to provide and maintain an **acceptable** level of service in the face of various faults and challenges to **normal** operation,

- in **Industrial** and **Organizational Safety**: …….. **Resilience Engineering** looks for ways to enhance the ability of organizations to create **processes** that are **robust yet flexible**, to monitor and revise risk models, and to **use resources proactively** in the face of disruptions or ongoing production and economic pressures. …… Success has been ascribed to the ability of groups, individuals, and organizations to **anticipate the changing shape of risk** before damage occurs; failure is simply the temporary or permanent absence of that.

* from Wikipedia and the book by Hollnagel, E., Woods, D. D. & Leveson, N. G. 2006. "Resilience engineering: Concepts and precepts", Aldershot, UK, Ashgate.

Questions to the Panelists:

- **How are resilience and resilience engineering approached in your Projects ?**
- **What methods and techniques are you investigating for obtaining resilient socio-technical complex systems ?**

DESEREC
Dependability and Security by Enhanced Reconfigurability

An *ICT for Trust and Security* research project
addressing
the dependability of Information systems

DESEREC
Dependability & Security by Enhanced Reconfigurability

Information Society Technologies

Dependability concerns

- The everyday life of European citizens relies on critical activities supported by networked Information Systems (I.S.):
  - **Communications** (telephone, Internet)
  - **Energy & fluids** (electricity, gas, water)
  - **Transportation** (railways, airlines, road)
  - **Health and emergency response**
  - **e-Government**

- So far, limited taken actions let these I.S.
  - not failure-proof enough to face:
    - **Software & hardware faults**
    - **Malicious actions: intrusion, virus**
  - with poor self-healing capability
    - **and therefore sensitive to cascading effects**
  - suffering long recovery time

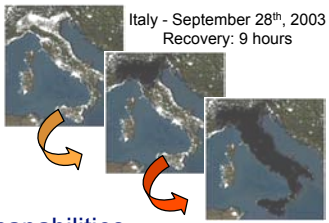

Italy - September 28th, 2003
Recovery: 9 hours

- The DESEREC project aims to leverage those capabilities
  - in new and existing Information Systems

Why DESEREC?

The picture

- Administrators are swamped by information of inappropriate level
- Most of the decision are taken short-term, with poor mid-term capability to arbitrate between business services with different criticality
- No synthetic view on dependability is provided

The proposed approach

- Provide information and interaction at *service* level instead of *component* level for day-to-day management
- Bring high-level management capabilities giving the ability to react appropriately upon errors/failures to maintain critical services
- Support mid-term strategy with planning and simulation tools enabling a proactive management of performance and dependability

The 3-tiered approach proposed by DESEREC

First objective – ***Detect & Prevent***

- Detect proactively incident and potential fault
- Keep as much as possible every failure local
  - Contain the incident: isolate the compromised area

Containment

Second objective - ***React***

- Sustain or quickly resume the critical applications
- Reallocate resources used by less critical ones

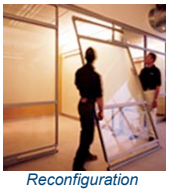

Reconfiguration

Third objective – ***Plan***

- Reallocate optimally the resources to recover the full range of services
- Validate the configurations by simulation

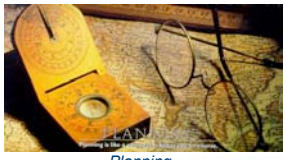

Planning

DESEREC - A multi-tiered response

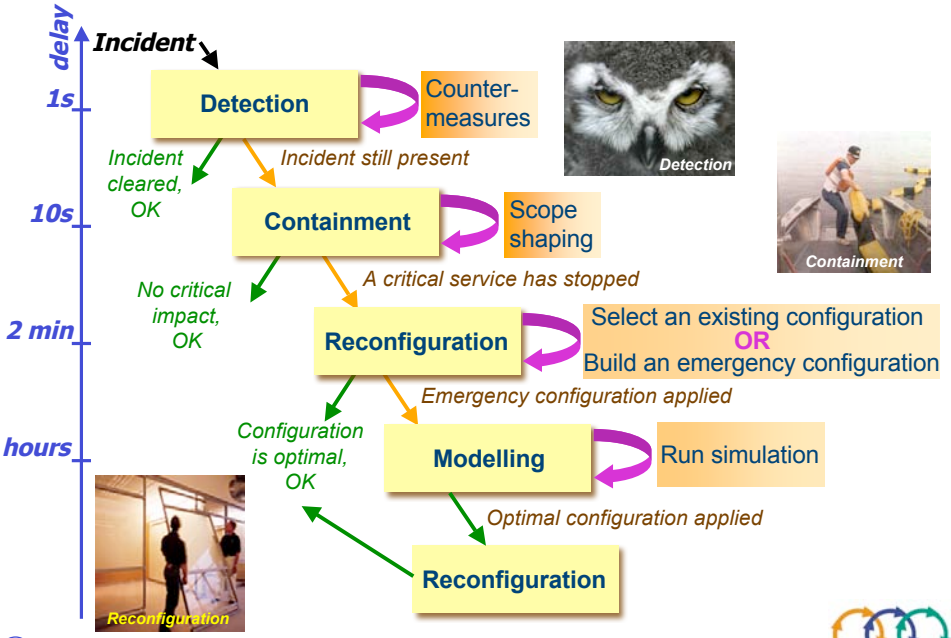

High level functional blocks

Management of mission-critical CIS via a model-based solution organised around three-tier reaction loop

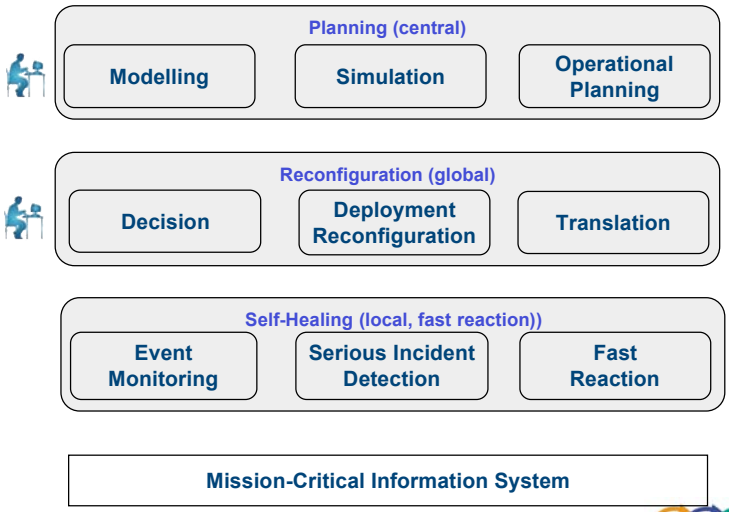

CIS seen as a cluster of molecules

- Introduce the molecule and multiple functional plans/views

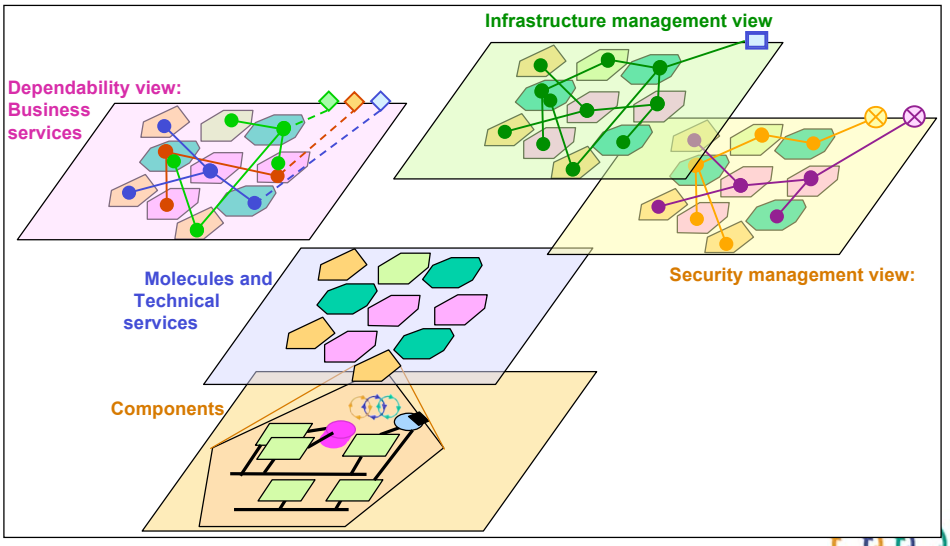

DESEREC Approach to resilience

- Optimizing the resilience of the Information System at the business service level
- The improvement of the resilience is achieved by optimizing the use of the available resources through reconfiguration
- Resilience engineering is one of the objective of DESEREC providing a learning mechanism for improving proactive reaction to incidents

ESFORS
European Security Forum for Web Services, Software and Systems

Panel on Resilience Views from other Projects

Presented by: Aljosa Pasic

Email: aljosa.pasic@atosorigin.com

Information Society Technologies

Funded by EC contract FP6-027599

Introduction

ESFORS

- **European Security Forum for Web Services, ESFORS**
- **European Technology Platform: Networked European Software & service Initiative , NESSI**

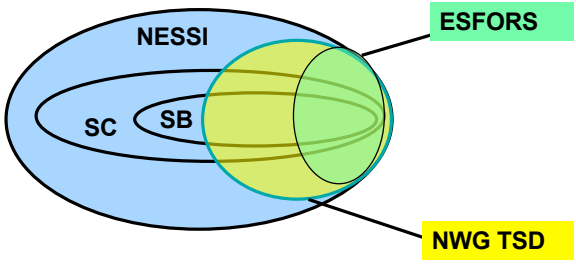

Atos Origin

hp invent

ENGINEERING INGEGNERIA INFORMATICA

SAP

WIT

FACULDADE DE CIÊNCIAS

Information Society Technologies

Funded by EC contract FP6-027599

ESFORS and Resilience

- **Applications will need to utilise shared and co-owned services out of different domains of control that require to obey separate security policies and ask for diverse security and dependability qualities**
- **What makes WS security different from other software components: trust is a driver for security requirements, accountability is a must**

Funded by EC contract FP6-027599

ESFORS and Resilience

- Driver 1: business resilience -> ICT resilience -> service and SOA resilience
- Driver 2: business functions decoupling -> software functions decoupling -> new dependencies (complex, dynamic, contextual) -> service/SOA resilience
- Driver 3: context information generated by SOA -> adaptability -> resilience

Funded by EC contract FP6-027599

ESFORS and Resilience

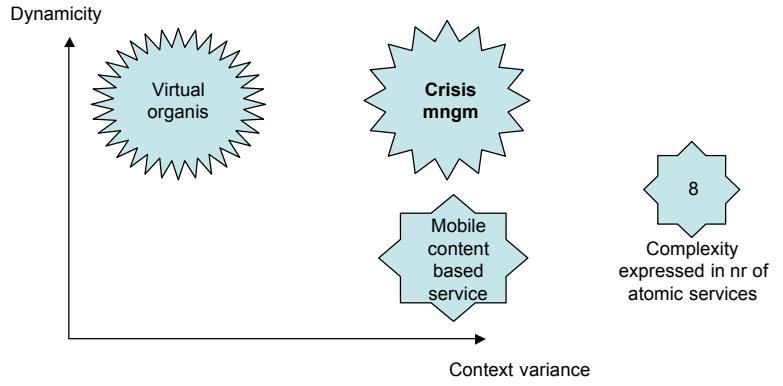

Funded by EC contract FP6-027599

ESFORS and Resilience

- Resilience in services vs resilience in SOA
- Resilience and concilience
- Resilience engineering throughout service lifecycle
- Resilience measuring

Funded by EC contract FP6-027599

ESFORS and Resilience

Dates to book: 10-11th of July 2007 !!!!!
Place: Maribor, Slovenia

Trust, security and dependability in service oriented applications and infrastructures: ESFORS workshop co-organised with NESSI, NESSI-Slovenia, Deserec, Serenity and Resist

Funded by EC contract FP6-027599

For more information:

Thank you

Aljosa Pasic (Atos Origin)
aljosa.pasic@atosorigin.com

Funded by EC contract FP6-027599

Domenico Presenza (Engineering)

Resilient SERENITY

Using S&D Patterns to enhance resilience

Budapest - 22/03/07

The SERENITY Objective

- To provide Security and Dependability (S&D) in Ambient Intelligence (AmI) scenarios.

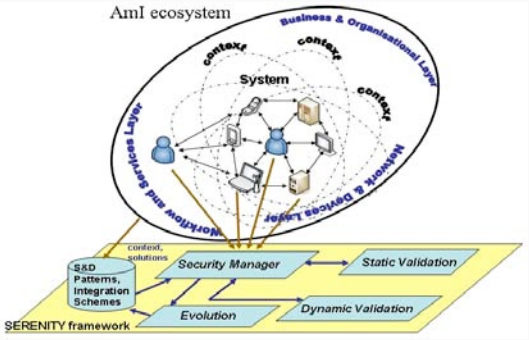

SERENITY Main Assumptions

- Security and Dependability knowledge can be coded (made explicit) through Security & Dependability Patterns (S&D Patterns);
- S&D Patterns can be integrated by means of Integration Schemes (IS);
- S&D Patterns can be monitored and, to some extent, enforced at run-time.

Aspects of S&D Solutions

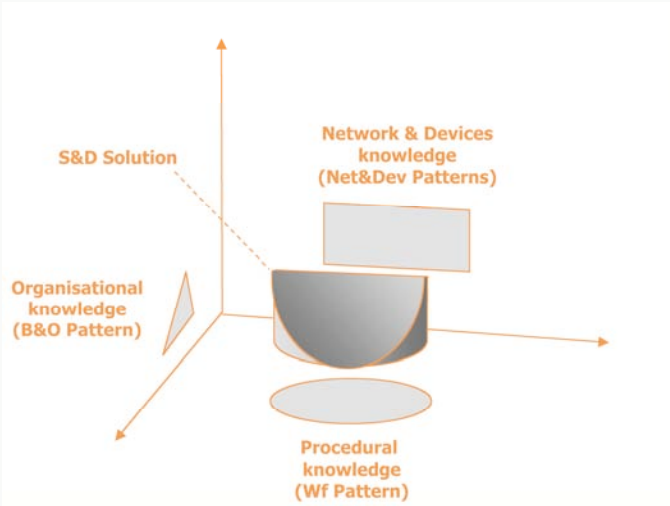

SERENITY Run-time Framework

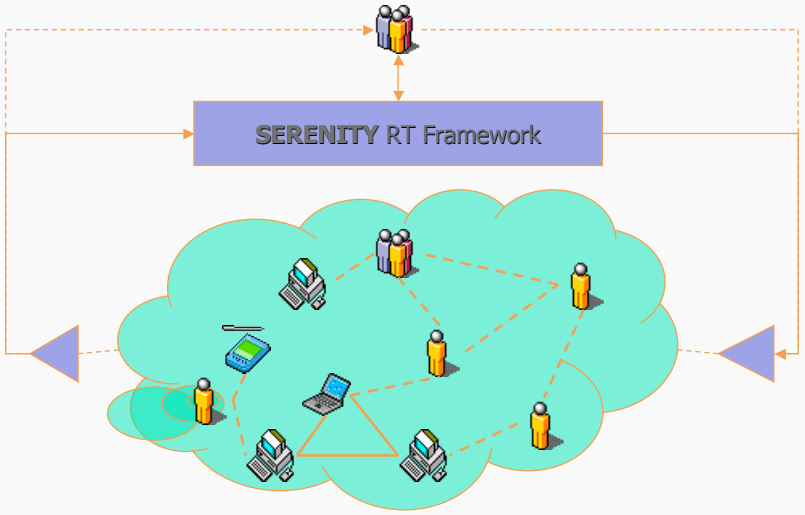

SERENITY about Resilience

- SERENITY is investigating whether S&D Patterns and Integration Schemes can be a tool to enhance the resilience of organisations by supporting run-time contextualisation of S&D management processes to the current situation.

S&D Patterns as RT Models

- An S&D Pattern/IS is used as an explicit representation of some portion of S&D Solution as perceived by Designers/Developers;
- Actors involved in S&D management interpret and adjust their behaviour to it;
- Prescribed part of the model are automatically interpreted and ambiguous/underspecified parts are left to the users for local adaptation, with tool support.

SERENITY Framework support

- The SERENITY RT framework will provide mechanisms for:
  - Monitoring integrity of S&D Patterns/IS and detecting deviations;
  - Monitoring status (QoS) of resources (services) available in the system;
  - Support coordination and communication of actors involved in the S&D management process;
  - Support on-line amendment of S&D Patterns/IS to reduce risks or cope with unexpected situations/threats.

HIDENETS – FP6 STREP

Scenarios and Resilience Solutions

Hans-Peter Schwefel, Aalborg University

hps@kom.aau.dk

HIDENETS Goals

- Develop and analyze **end-to-end resilience solutions**
  - for scalable distributed applications and mobility aware services
  - in ubiquitous communication scenarios
    - Example use-case: car2car communication with server-based infrastructure
  - assuming highly dynamic, unreliable communication infrastructures
- Planned **results** are
  - architectural and design solutions
  - communication protocol extensions and dependability middleware
  - methods for quantitative analysis and testing
  - tools for development and analysis

  for end-to-end system level resilience and dependability
    - based on standard off-the-shelf components
    - in wireless communication networks and infrastructure-based settings

HIDENETS Scenarios

- **Applications** with varying dependability requirements, e.g.
  - Platooning
  - Floating car data, hazard warning
  - Distributed black-box
  - Streaming (video/data)

- **Challenges** of the C2C/C2I scenarios
  - Dynamicity/Mobility: changing topologies and communication characteristics in ad-hoc domain and in connection to infra-structure services
  - Open systems with (C)OTS components
  - Heterogeneity: different network domains [and different node capabilities]
  - Resource limitations and strong cross-influence between system parts
  - Accidental and malicious faults

  + large number of nodes, privacy aspects…

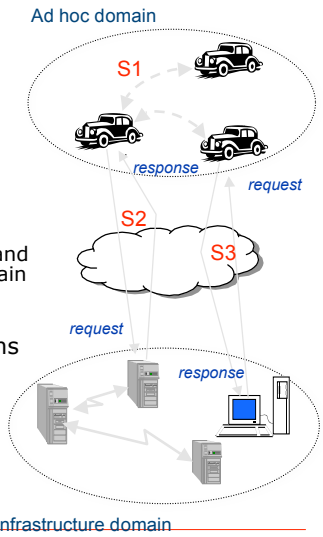

HIDENETS Approach

- Steps (inter-linked)
  - Applications/use-cases → requirements → necessary middleware and communication layer functions
  - Network and node architecture → fault-models → detailed function/algorithm/protocol development, experimental implementation, modeling and assessment

- Resilience solutions: joint optimization via
  - Differentiation
    - Architectural: wormhole concept
    - Flow/packet/message treatment: scheduling/routing/etc.
  - Fault detection and recovery, as well as masking
    - Communication interfaces/links/paths: interface selection, (multi-path) routing, Gateway selection
    - Node functions: data storage, computations
  - APIs that allow for adaptive applications

  While maintaining the end-to-end, holistic system view, covering
    - All nodes on the end-to-end path
    - Communication protocols as well as service middleware

HIDENETS node software architecture

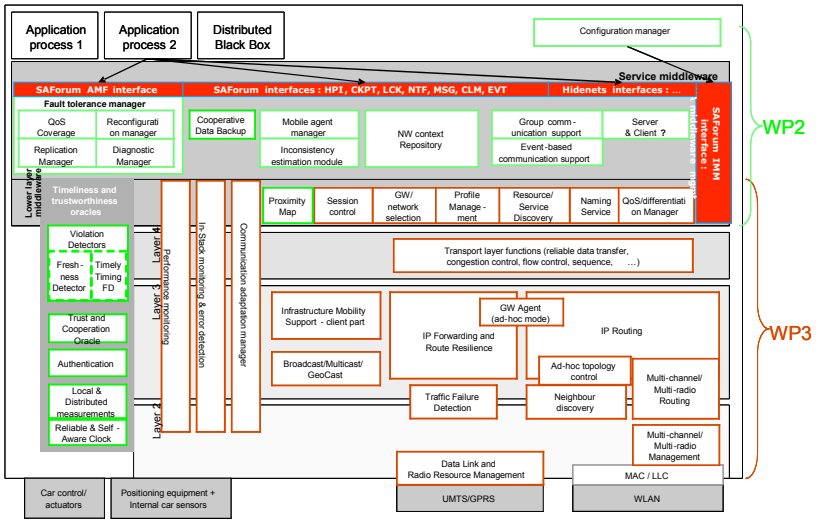

HIDENETS hybrid architecture

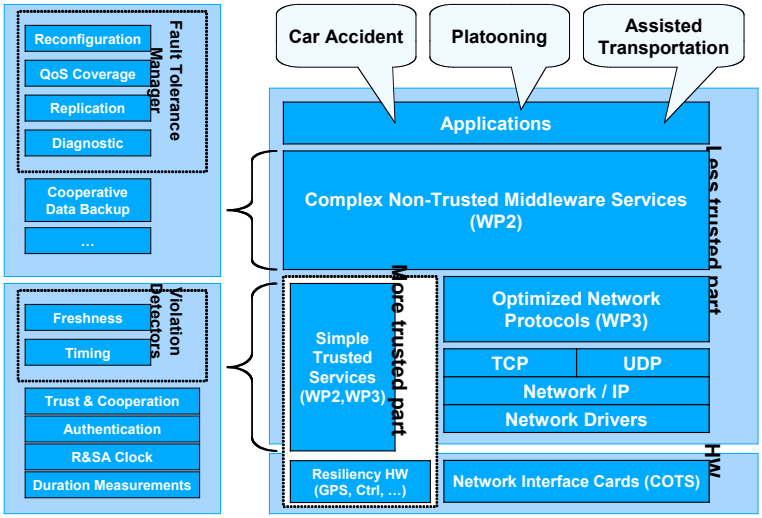

Summary

- Goal: end-to-end resilience solutions for car-to-car and car-to-infrastructure scenarios
  - Communication protocols (L2-L4), middleware functions, application interfaces, application development tools
  - Mainly (but not exclusively) accidental faults: communication links and nodes (both in ad-hoc and infrastructure domain)
  - Interaction of resilience mechanisms while still keeping a layered structure
  - Assessment in analytic/simulation models, and experimental set-ups

Technical deliverables are available on web-page: www.hidenets.aau.dk
Final results: December 2008

ReSIST

Resilience for Survivability in IST

A European Network of Excellence

- Auto-evaluation
- The steps forward

☞ Objectives

1) *Integration* of teams of researchers so that the fundamental topics concerning scalably resilient ubiquitous systems are addressed by *a critical mass* of co-operative, multi-disciplinary research
2) Identification, in an international context, of the key *research directions (both technical and socio-technical)* induced on the supporting ubiquitous systems by the requirement for trust and confidence in ambient intelligence
3) Production of significant *research results (concepts, models, policies, algorithms, mechanisms)* that pave the way for scalably resilient ubiquitous systems
4) Promotion and propagation of a *resilience culture* in university curricula and in engineering best practices

☞ Coverage of the objectives

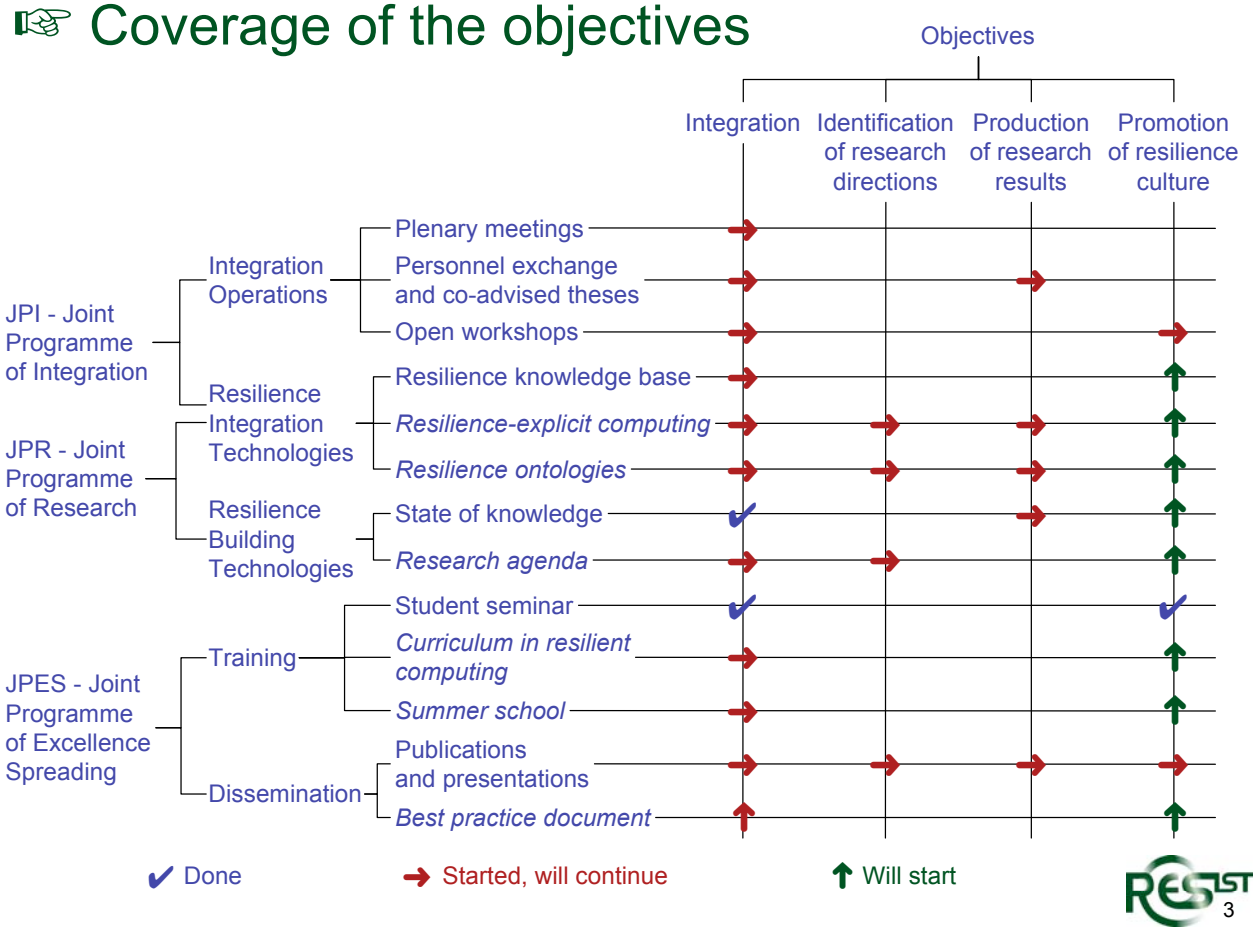

☞ First year

Intense work …

- For the record
  - State of knowledge document
  - Prototype knowledge base
  - Preparatory ground work

… supported by numerous meetings …

- 1 plenary meeting
- 5 executive board meetings
- 6 committee (RKB, T&D) meetings
- 1 student seminar
- 1 SIG (resilience ontology) meeting
- 4 WG (2 Socio, 1 Arch-Algo, 1 Verif) meetings

… and by a Wiki …

… that gave impetus to integration and community building

☞ Second year

Continuation of intense work …

- For the record
  - Research agenda according to the resilience-scaling technologies
    - Evolvability
    - Assessability
    - Usability
    - Diversity
  - Support for resilience-explicit computing first edition
  - Resilience knowledge base version 2
  - Resilience ontology
  - Resilient computing curriculum draft
  - Resilient computing courseware outline
  - Summer school
  - Best practice document outline

… open to external contributions …

- Already planned actions
  - Critique of the research agenda
  - Establishement of resilient computing curriculum
  - Definition and production of the best practice document
- Creation of *affiliate* status

… supported by an overhauled website

- Contents and design