Socio-Technical Complex Systems of Systems: Can We Justifiably Trust Their Resilience?

Luca Simoncini

University of Pisa, Department of Information Engineering, Via G. Caruso 16, 56122 Pisa, Italy

e-mail: <luca.simoncini@iet.unipi.it>

Complex systems need to be resilient

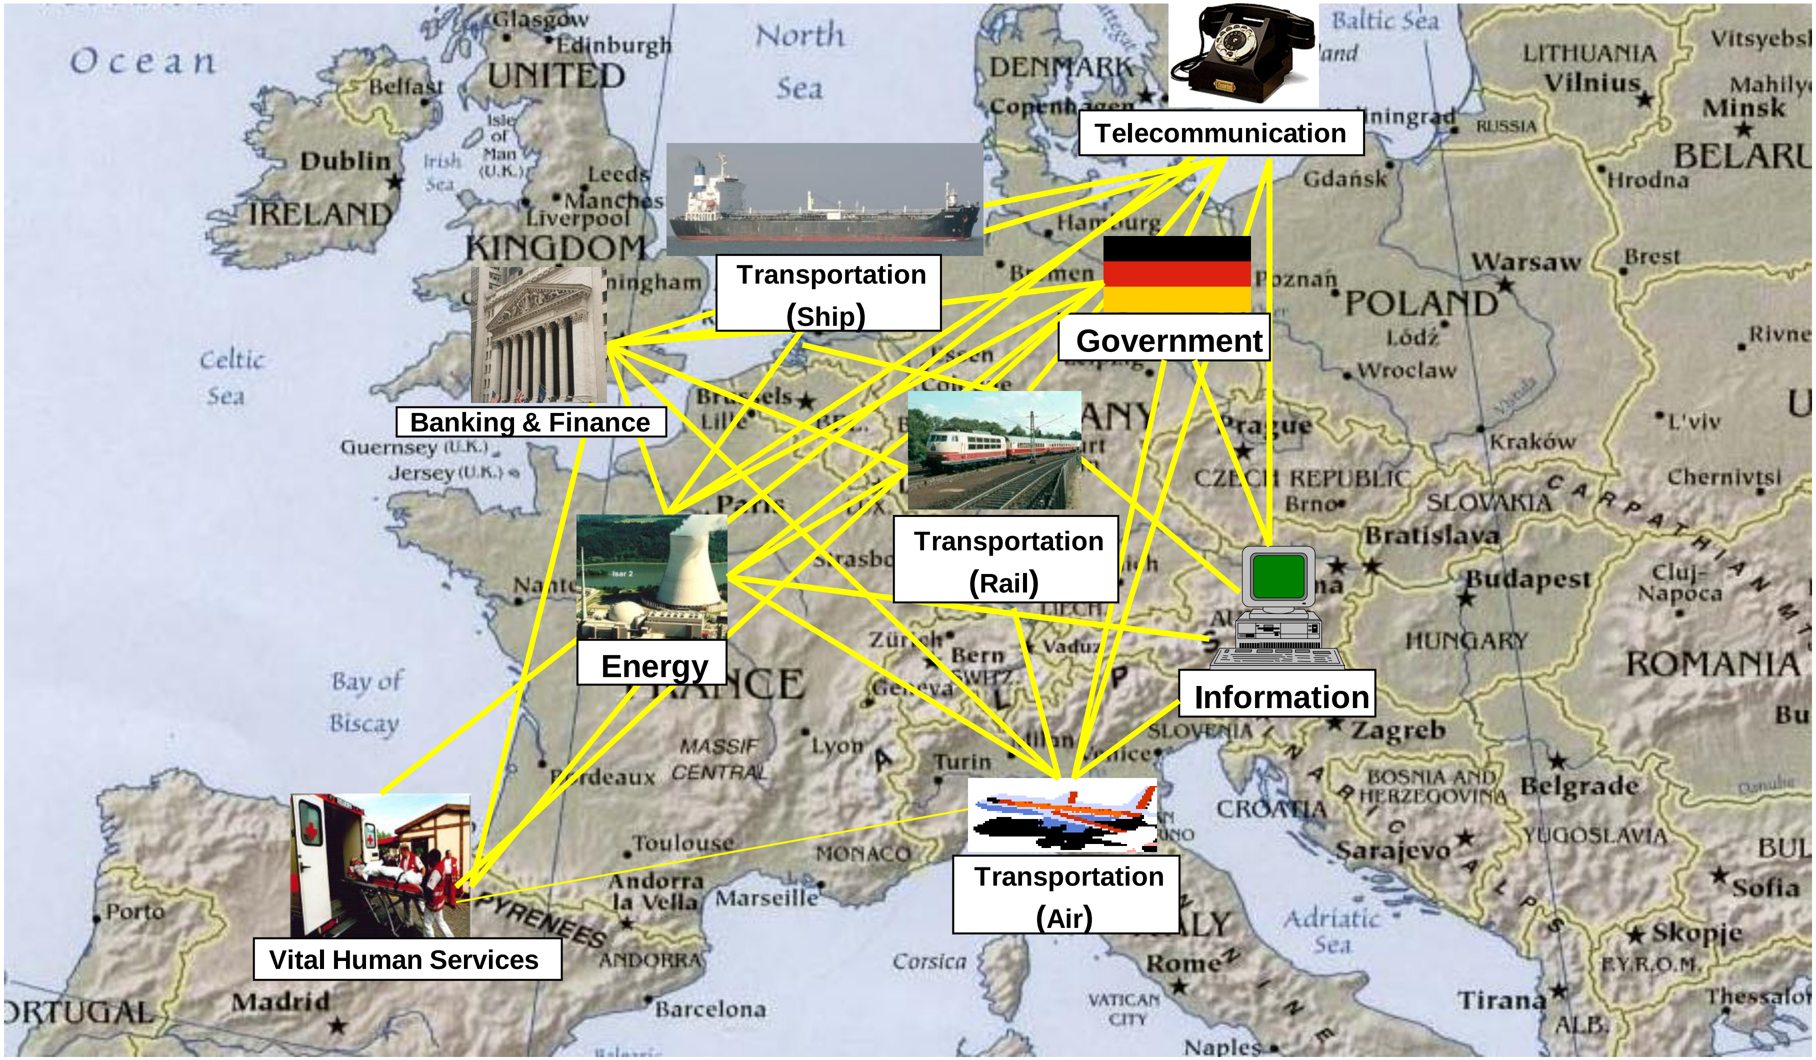

Telecommunication

Transportation (Ship)

Government

Banking & Finance

Transportation (Rail)

Energy

Information

Transportation (Air)

Vital Human Services

On the term Resilience

The term resilience has been used in many fields and, as a property, two threads can be identified: a) in social psychology, where it is about elasticity, spirit, resource and good mood, and b) and in material science, where it is about robustness and elasticity.

The notion of resilience has then been elaborated:

- **In child psychology and psychiatry, referring to living and developing successfully when facing adversity;**
- **In ecology, referring to moving from a stability domain to another one under the influence of disturbances;**
- **In business, referring to the capacity to reinvent a business model before circumstances force to;**
- **In industrial safety, referring to anticipating risk changes before damage occurrence.**

A common point to the above senses of the notion of resilience is the ability to successfully accommodate unforeseen environmental perturbations or disturbances

Resilient Computing and Resilience Engineering

Resilience (for computing systems and information infrastructures):

the persistence of service delivery that can justifiably be trusted, when facing changes

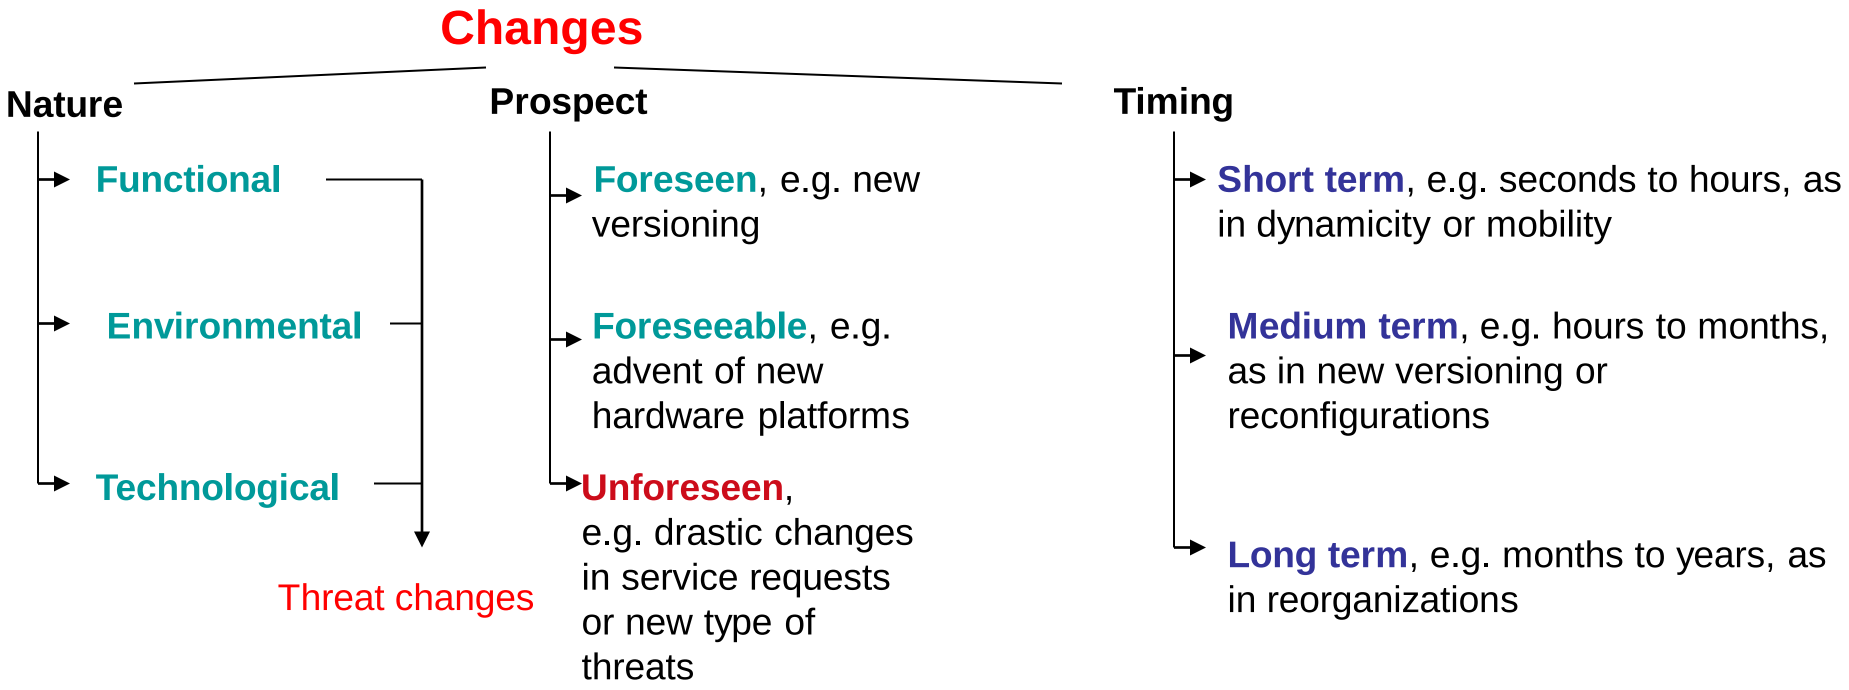

Changes

Nature

- **Functional**
- **Environmental**
- **Technological**

Threat changes

Prospect

- **Foreseen**, e.g. new versioning
- **Foreseeable**, e.g. advent of new hardware platforms
- **Unforeseen**, e.g. drastic changes in service requests or new type of threats

Timing

- **Short term**, e.g. seconds to hours, as in dynamicity or mobility
- **Medium term**, e.g. hours to months, as in new versioning or reconfigurations
- **Long term**, e.g. months to years, as in reorganizations

Resilience Engineering:

How to design, implement, operate etc. complex systems so that they can be resilient

What about resilience in Fukushima Dai-ichi accident ?

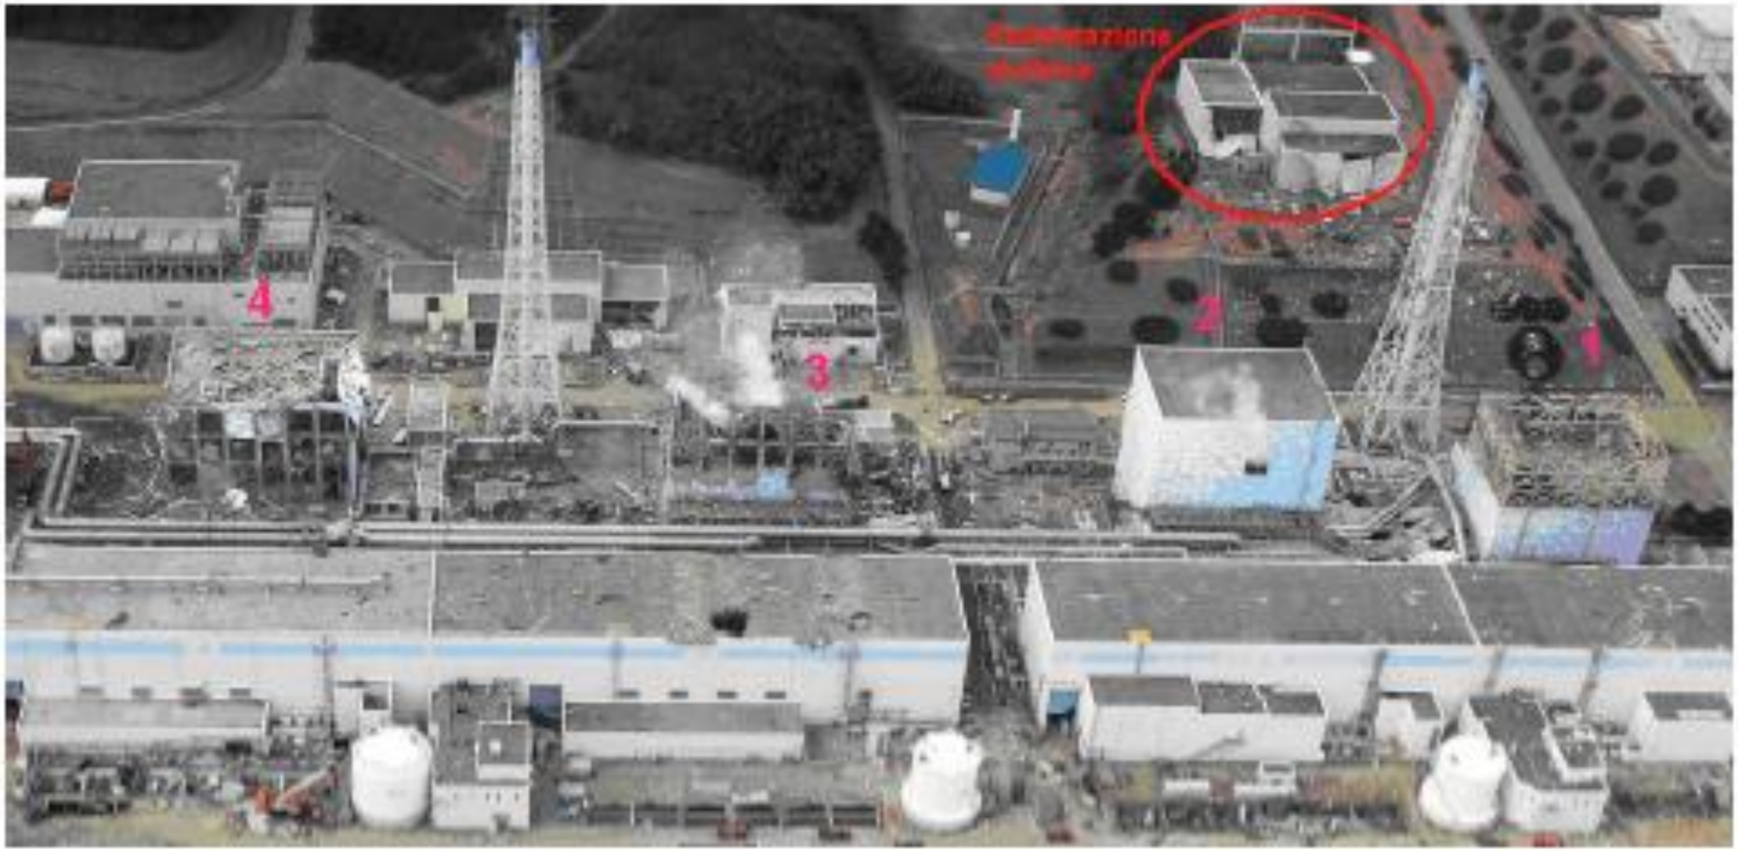

Schematic interdependent blocks in Fukushima Dai-ichi

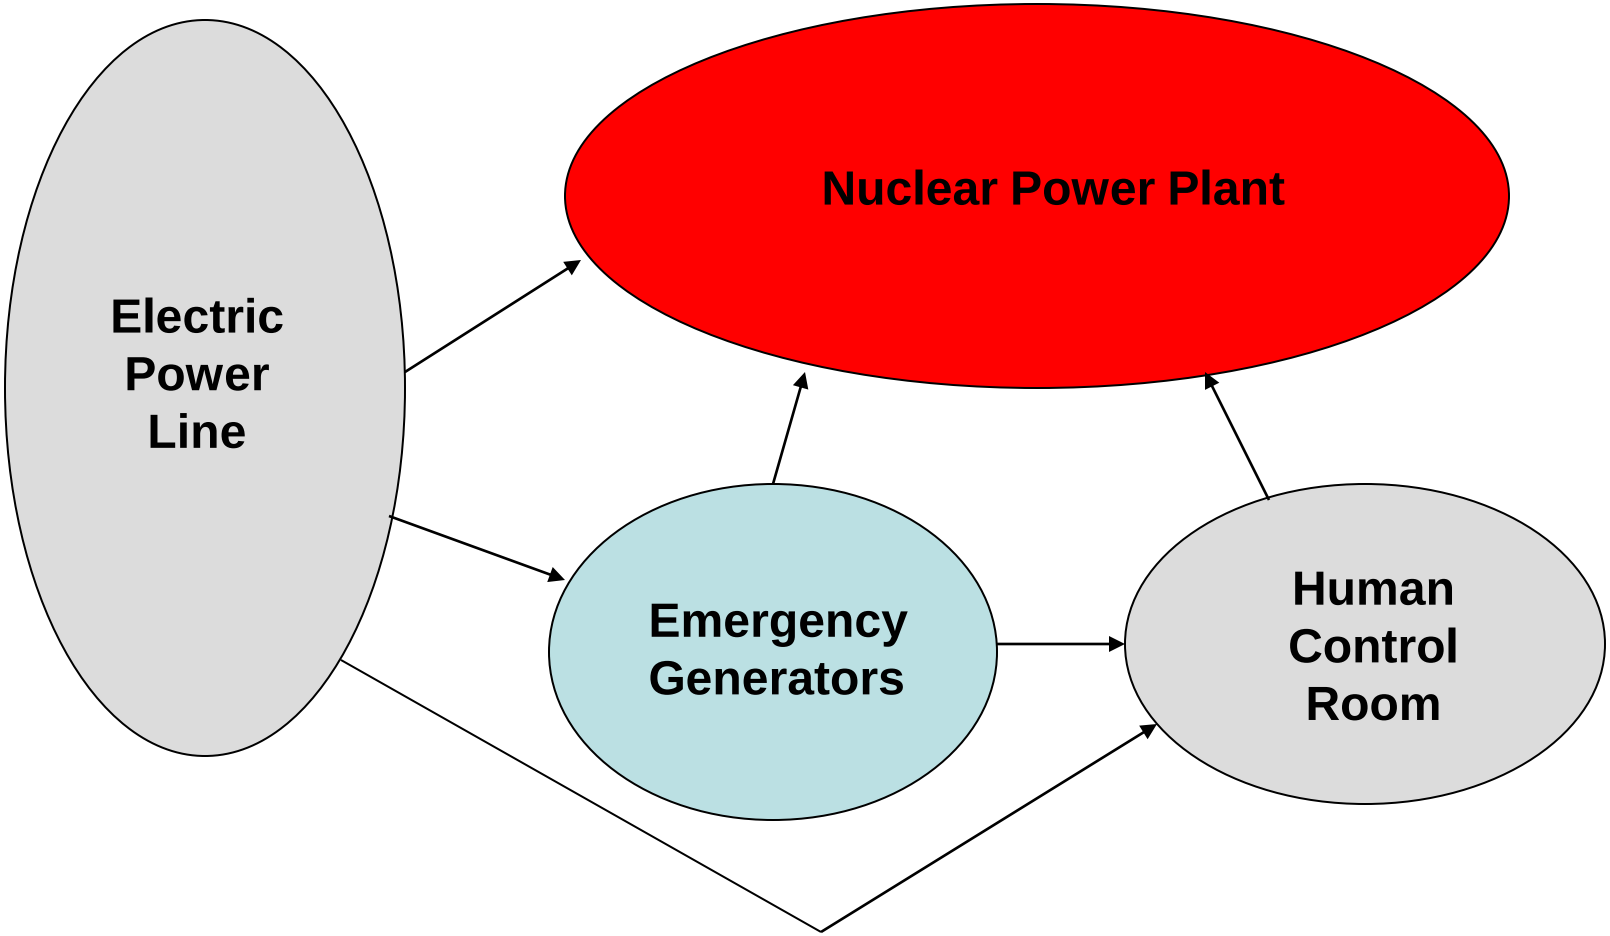

Electric Power Line

Nuclear Power Plant

Emergency Generators

Human Control Room

Section of Fukushima Dia-ichi nuclear power plant

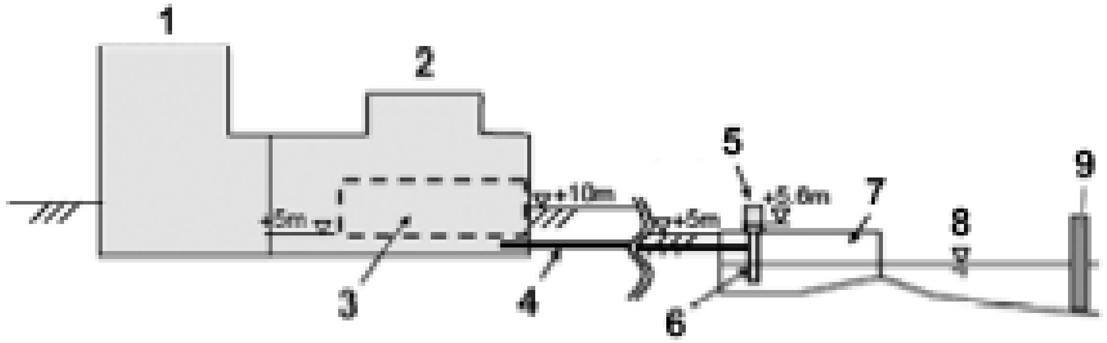

1 Reactor Building

2 Turbine building

3 Emergency diesel generators room

4 Sea water pipe

5 Engine for sea cooling water

6 Pump for sea cooling water

7 Sea water pool

8 Sea level

9 Sea barrier (10 meters/33 feet high)

Container Mark1 of Units 1, 2, 3, 4, 5 of Fukushima Dai-ichi

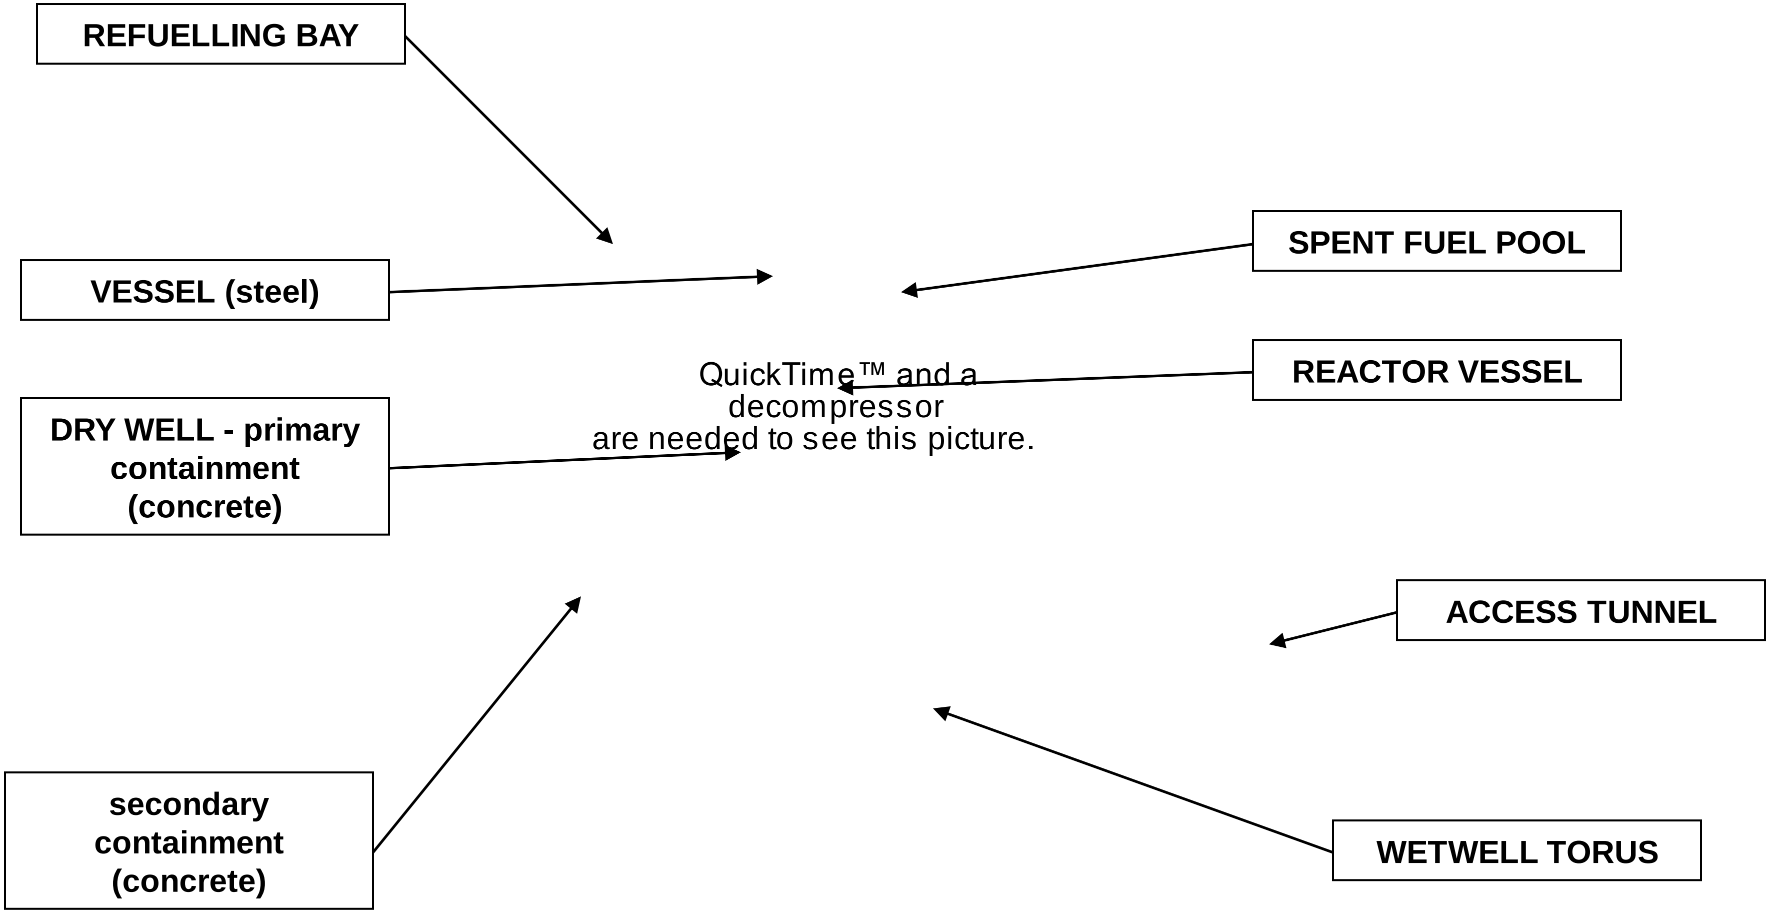

REFUELLING BAY

VESSEL (steel)

DRY WELL - primary containment (concrete)

secondary containment (concrete)

SPENT FUEL POOL

REACTOR VESSEL

ACCESS TUNNEL

WETWELL TORUS

QuickTime™ and a
decompressor
are needed to see this picture.

Sequence of events March 11, 2011

14:46 Heartquake Magn.9 (Magn. 7 at Fukushima - duration 3 minutes) - Consequence: total electric power interruption (6 external lines + the all internal switching power sub-stations)

15:30 Tsunami with waves 14/15 meters - 46/50 feet high - Consequences: complete flooding of the nuclear power plant. Complete destruction of the main internal electric power sub-station. All area - internal and neighbour roads - interrupted and covered with more than 5 meters mud and debris. Destruction of sea water cooling systems, turbine buildings flooded with destruction of emergency diesel generators, control rooms out of work, instruments (analogic) no more usable, no communication to the external, human reaction in extraordinary situation.

Even if

2 seconds after electric power interruption, anti-seismic emergency controls have turned-off all active nuclear reactors, main steam turbines have been bypassed and the main three safety functions (control of the nuclear process, heat control of the core and confinement of radioactive material) correctly started

anyway

the **combined effects** of heartquake and tsunami provoked the disaster

Weak Resilience of the System

1) External electric power interruption due to the earthquake:

- External electric supply is essential during an emergency - No dual redundant external supply provided

2) Sea barrier:

- High 10 meters/ 33 feet not sufficient to cope with tsunami's waves

2) Location of emergency diesel generators:

- Dual redundancy used **but** located in the same place (below street level) so exposed to common mode failures (tsunami)

3) Location of control rooms and analog instruments:

- No remote (digital) data acquisition or possibility of intervention from a safe location

4) **No (or very poor) communication and psychological effects on trained maintenance personnel:**

- Only partial knowledge local to each reactor. Very stressful situation

Resilience engineering not correctly used in design, implementation, off-line stress testing and identification of emergency protocols

B. Littlewood and ultra-high dependability

Littlewood, B.: Limits to Evaluation of Software Dependability. In: Software Reliability and Metrics, 7th Annual CSR Conference, Garmisch-Partenkirchen

“The first point I wanted to make was that there is an inherent **uncertainty** about the behaviour of the systems we build, which forces us to talk about their dependability in the language of probability and statistics. This is unpalatable to many computer scientists, who conclude from the deterministic predictability of the machine itself that we can have similar determinism at the macroscopic level of the behaviour observed by a user. …..

If readers concede this first point, it seems to me that questions about whether a software-based system is fit for its purpose become **questions about whether a particular probabilistic dependability level has been achieved** - a problem of numerical evaluation. It is easy to show that we can only answer such questions when the level required is quite modest: we ***cannot*** gain confidence in ultra-high dependability **without** obtaining (literally) **incredible amounts of evidence**. …..

I think these difficulties have serious implications for the builders of safety-critical systems, and for society at large. It is easy to be seduced by the extensive functionality that can be provided by software, without the constraints of ensuing hardware unreliability. Some of the benefits from this functionality may indeed be claimed to bring enhanced safety. But at the end of the day **we have a right to demand that the system is sufficiently safe, and this cannot be demonstrated for some of the systems that we are building even now. Perhaps now is a time to take stock and consider some retrenchment - for example, deciding that an unstable civil airliner is not a good thing**.

All is not gloom. Systems with these dramatic requirements are not that common. For more modest systems, with modest dependability requirements, evaluation is possible. Further research will extend this achievement - but only relatively modestly.

Complex Systems of Systems and Resilience

Fukushima plant and services are an example of a composition of “mature technology” socio-technical systems operated by highly trained personnel, designed under very consolidated safety standards that anyway produced a very severe disaster for people and environment.

If we extend our attention to present market trends that push forward the deployment and operation of **complex software-based socio-technical systems of systems** that have to be not only dependable but also resilient,

the following question becomes mandatory

Can we justifiably trust the resilience of such systems ?

Peculiarities of Socio-Technical Complex Systems of Systems

- These systems have to be dependable in a highly dynamic and changing environment, that is they have to present a high level of resilience;
- These systems are always on line and require pervasive and ubiquitous computing;
- They have to manage sensitive personal, social, corporate, financial and political information;
- These systems include a huge number of humans, as users or operators, who are often untrained and often risk unaware.

These systems have to cope with **new fault types** (escalating and cascading failures in critical utility infrastructures, identity theft, unexpected interactions, common mode failures affecting a very large number of users, etc.), with **changing environments** and possible **uncontrollable human interactions**.

Open research topics

- Understanding the new risks and threats;
- Understanding the boundary-less nature of systems;
- Dealing with increased scale and complexity and criticality;
- An assessment based on user perception
- Dealing with changing environments.

but

it is impossible to anticipate all the possible situations and events that could happen and that could lead to failures with possible catastrophic consequences. This means that we are going to operate quite critical systems whose design has been made in ***ignorance*** or in ***complete unawareness*** of their requirements.

It is evident that a correct and accurate assessment of the resilience of these systems is questionable or impossible.

What about requirements ?

The requirements for a system can be clustered into four groups:

- the **Known Knowns** – what we know that we know
- the **Known Unknowns** – what we know that we do not know
- the **Unknown Knowns** – what we pretend not to know even if we know
- the **Unknown Unknowns** – what we do not even know that we do not know

The KK and KU groups are the easiest since they include all requirements that can be deterministically considered in the design

The specifications for all functions of the system may be derived.

We can also take into consideration (maybe rare) events that we may not be able to deterministically express but that we know that could happen even if we do not know either when they will happen or how they will manifest themselves (some type of security attacks)

What about UK and UU groups?

The UK group is intriguing since it refers to things that we know may (or will) happen but we pretend not to know.

This is quite common in social and political affairs, when “we close our eyes to the evidence” and expect to see what happens.

In our framework, maliciously or accidentally neglected specifications belong to this group, which then lead to incorrect designs, or erroneous operations performed under stress or time pressures.

The UU group is clearly the most dangerous since there is no way of being able to consider something that we are completely ignorant - or unaware of - in terms of its possibility, manifestation (what, when and how) and consequences

The type and number (that cannot be known) of possible events belonging to UU group (by experience) grow according to the complexity of the system and its interactions. Their manifestation and consequences cannot be known or forecast. This is a very vague statement that may be mitigated by qualifying these events as “**extremely rare and unlikely**”, however there is a **contradiction** since this qualification applies to events of which we have some knowledge

Black Swan Events - 1

The theory of Black Swan Events (introduced in 2007 by Nassim N. Taleb) was developed to explain:

- The disproportionate role of high-impact, hard to predict, and rare events that are beyond the realm of normal expectations in history, science, finance and technology
- The non-computability of the probability of the consequential rare events using scientific methods (owing to their very nature of small probabilities);
- The psychological biases that make people individually and collectively blind to uncertainty and unaware of the massive role of the rare event in historical affairs

From Taleb: “What we call here a Black Swan (and capitalize it) is an event with the following three attributes. First, it is an **outlier**, as it lies outside the realm of regular expectations, because nothing in the past can convincingly point to its possibility. Second, it carries an **extreme impact**. Third, in spite of its outlier status, human nature makes us concoct explanations for its occurrence after the fact, making it **explainable and predictable**. I stop and summarize the triplet: **rarity**, **extreme impact**, and **retrospective (though not prospective) predictability**”

Black Swan Events - 2

- The event is a surprise (to the observer);
- The event has a major impact;
- After the fact, the event is rationalized by hindsight, as if it *could* have been expected (e.g. the relevant data were available but not accounted for)

The main idea for coping with Black Swan Events is **not to attempt to predict them**, but **to build robustness** against negative ones that occur and be able to exploit positive ones by identifying areas of vulnerability in order to **"turn the Black Swans white"**

Black Swan Events are phenomena with specific empirical and statistical properties, which Taleb puts in "the fourth quadrant"**.**

The problem concerns the decrease in knowledge when it comes to rare events as these are not visible in past samples and therefore require a strong a priori, or what one can call an extrapolating theory. Accordingly events depend more and more on theories when their probability is small. In the fourth quadrant, knowledge is both uncertain and consequences are large, requiring more robustness.

Black Swan Events - 3

Before Taleb, those who dealt with the notion of the improbable, such as Hume, Mill, and Popper focused on the problem of induction in logic, specifically, that of **drawing general conclusions from specific observations**. Taleb's Black Swan Event has a central and unique attribute, ***high impact***. His claim is that ***almost all consequential events in history come from the unexpected*** - yet humans later convince themselves that these events are explainable in hindsight (bias)

One problem is the belief that the unstructured randomness found in life resembles the structured randomness found in games. This stems from the assumption that the unexpected may be predicted by extrapolating from variations in statistics based on past observations, especially when these statistics are presumed to represent samples from a bell-shaped curve

More generally, decision theory, which is based on a fixed universe or a model of possible outcomes, ignores and minimizes the effects of events that are "outside the model". **A fixed model considers the "known unknowns", but ignores the "unknown unknowns"**

Consequences

When complex decisions need to be taken in the fourth quadrant, **neither statistics nor models can be used**, and worse, if they are used following a classical approach, they will deceive us.

Restricting the horizon to the design of socio-technical complex systems of systems in changing environments the application of the Black Swan Events theory implies the **impossibility of providing a meaningful and convincing assessment of their resilience based on classical statistical methods**

Thus the main question must be changed from "How can resilience be quantitatively assessed?" to **"How can a socio-technical complex system of systems in changing environments be designed in order to increase the confidence that it may survive and provide an acceptable (maybe reduced) service not only in the presence of changes but also of Black Swan Events (unknown unknowns)?"**

Primary concerns

- Due to unknown unknowns, failures will happen and we cannot forecast their evolution (whether or not there will be catastrophes) or probability distribution
- It is impossible to have complete and accurate specifications of such systems, therefore capturing requirements is an evolving process not only during the design but also during the operational life
- A socio-technical complex system of systems is not built from scratch, but is the composition of pre-existing systems. It is based on their individual infrastructures, each designed with different paradigms, with different forms of governance, and made to cooperate through interfaces (from input signals to communicating humans)
- Humans (designers, maintenance personnel, operators and end-users) are significant components of these systems. It is clear that it is impossible to force requirements on humans, especially if the system is used and operated by a large number of persons who are untrained and/or risk unaware
- The stored and managed information are sensitive from many points of view

What can be done ?

We need to approach the design, deployment and operation with tools and methods that stem from engineering (computer, software, ergonomics etc.), decision-making and management, social sciences, culture and education in a holistic way

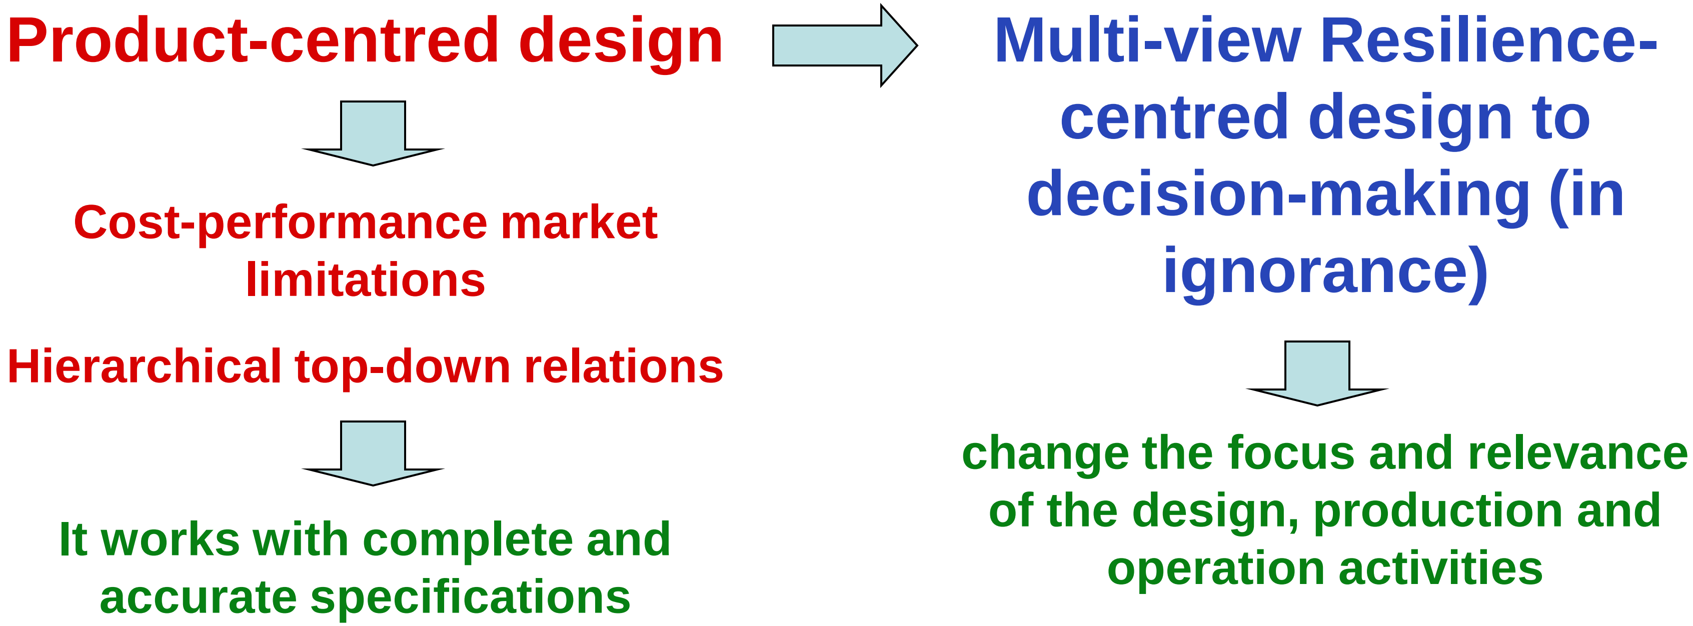

Product-centred design

Cost-performance market limitations

Hierarchical top-down relations

It works with complete and accurate specifications

Multi-view Resilience-centred design to decision-making (in ignorance)

change the focus and relevance of the design, production and operation activities

What can be done ? - Architectural capabilities

The most important requirements in making decisions in ignorance are:

- Detect the incorrectness of the decision as soon as possible. It requires the enhancement of the ***controllability*** of the design process and of the designed system
- Correct the decisions. It requires the ***flexibility*** and ***adaptability*** of the design process and of the designed system

Controllability, ***flexibility*** and ***adaptability*** are introduced as three capabilities that need to be exploited in socio-technical complex systems of systems

They are architectural capabilities and require a **special focus in the design and operation of the interfaces**. In fact a system may fail because of the interactions between components that operate as designed

All interfaces are the weak point in all systems of systems

(signals and commands, human communication during design, operation, maintenance, warning and emergency management,….)

Less fully automated procedures ⇒ **Main human responsibility + automated early warning detection**

What can be done ? - Technical level

- From B. Randell (2005): “However, for structuring to have some direct relevance to questions of operational dependability, and in particular fault tolerance, it must be what might be described as strong – **strong** structuring actually controls interactions within and between systems, and limits error propagation in both time and space, i.e. constitutes *real* not just perceived or imagined boundaries.”

- From C. Jones (2007): a systematic way to derive specifications and leads to recording assumptions concerning the controlled world to specify and design complex systems of systems

- From M. Thomas (2003): poses the question, with respect to software dependability, as to what is ***known*** and what is ***not done*** in industry due to educational and cultural limitations and difficulties deriving from cost-performance pressures.

- Recent EU Projects and NoE:

- Crutial (2006-2008) - networked ICT systems for the management of the electric power grid
- ReSIST NoE (2006-2009) - Resilience for Survivability in IST
- Hidenets (2006-2009) - end-to-end resilience solutions for distributed applications
- CONNECT (2009-1012) - How to make networked systems eternally connected

What can be done ? - Managerial protocols

- Changing the focus from product-centred to resilience-centred design and operations requires a strong shift in the managerial skills and attitude
- Each decision-maker (at design and operational level) should have a complete view of the mechanisms and types of interactions between the cooperating systems
- Special training is needed aimed at a deep revision of the current cultural protocols in the managerial, technical and responsibility chains
- An independent authority, not tied to the protocols to be changed, is needed
- A new resilience-centred organization should execute technical and managerial processes as if they were a single process
- In case of “unknown unknowns”, a multi-view (controllability) and shared analysis of the manifestations (flexibility) should be undertaken rapidly in order to introduce possible countermeasures (adaptability) with no centralized and delegated responsibility

What can be done ? - Educational issue

What about the educational and cultural attitude of the multitude of untrained and risk unaware users who operate and interact with the system?

- Open door to the system with easy access for malicious activities (attacks to privacy, identity theft, malicious matching and profiling, to sabotage and terrorist attacks even those designed simply to cause widespread panic)
- Unpredictability of events generated by untrained users
- Evident social impact
- No way of solve this problem by only spreading education and culture on risks and consequences
- Toughest to address since it requires the spread of resilience-centred attitudes to a large patchy set of interacting entities, many of which have limited knowledge of their actions

What can be done ? - Conclusions

- Research on dependability attributes and means precondition for the ability to develop and operate resilient systems
- Composing dependable systems does not imply the dependability of the composed system and least of all its resilience, due to the impossibility of providing statistical measures, the changing environment and the presence of “unknown unknowns”
- Using the best known heuristics to design and operate socio-technical complex systems of systems in changing environments
- Put the maximal attention to the design and operation of interfaces
- Attention to small events and warnings that may indicate a deviation from expected behaviour, in order to try to build some qualitative metric for resilience

Anyway

Socio-technical complex systems of systems in changing environments will fail sooner or later, with very difficult forecast of the consequences of such failures

Introduction of socio-technical complex systems of systems and extension of civil and social rights for citizens against an excessive and dangerous pervasiveness

Suggested readings

- Simoncini, L.: Socio-technical Systems of Systems: Can We Justifiably Trust Their Resilience. In: Dependable and Historic Computing: Essays dedicated to Brian Randell on the Occasion of His 75th Birthday, Jones, C.B. and Loyd, J.L. (eds.), Lecture Notes in Computer Science, 6875, pp 486-497, Springer-Verlag (2011)
- Littlewood, B.: Limits to Evaluation of Software Dependability. In: Software Reliability and Metrics, 7th Annual CSR Conf., Garmisch-Partenkirchen, pp. 87-110, Elsevier (1991)
- Laprie, J.C.: From Dependability to Resilience. Technical report LAAS n. 08001 (2008)
- Taleb, N.N.: The Black Swan: The Impact of the Highly Improbable. Random House and Penguin. ISBN 978-1-4000-6351-2. New York (2007). The book was completed in 2010 with the second edition including a long essay "On Robustness and Fragility”
- Taleb, N.N.: The Role and Nature of High Impact Events (Black Swans): Technical Commentary and Empirical Data. (2008) http://www.fooledbyrandomness.com/EDGE/index.html
- Jackson, S.: Architecting Resilient Systems. John Wiley & Sons, Hoboken, NJ (2010)
- Randell, B.: Dependability, Structure and Infrastructure. In: Cyberspace Security and Defence: Research Issues. Kowalik, J.S., Gorski, J. and Sachenko, A. (eds.), pp. 143--160, NATO Science Series II, ISBN 1-4020-3380-X, Springer (2005)
- Jones, C.B., Hayes, I.J., Jackson, M.A.: Deriving Specifications for Systems that are Connected to the Physical World. In: Formal Methods and Hybrid Real-Time Systems: Essays in Honour of Dines Bjørner and Zhou Chaochen on the Occasion of Their 70th Birthdays. Jones, C.B., Liu, Z. and Woodcock, J. (eds.), Lecture Notes in Computer Science, 4700, pp 364-390, Springer-Verlag (2007)