INFOCOMP 2013

Exa-Intelligence:
Next Generations of Intelligent Multi-Agent and High End Computing Systems in Development and Practice

November 18, 2013, Lisbon, Portugal

The Third International Conference on Advanced Communications and Computation (INFOCOMP 2013)

INFOCOMP
November 17–22, 2013 - Lisbon, Portugal

INFOCOMP International Expert Panel: Exa-Intelligence

Panelists

- *Claus-Peter Rückemann* (Moderator), Leibniz Universität Hannover / Westfälische Wilhelms-Universität Münster (WWU) / North-German Supercomputing Alliance (HLRN), Germany
- *Paulo Leitão,* Polytechnic Institute of Bragança, Bragança, Portugal
- *Udo Inden,* Cologne University of Applied Sciences, Cologne, Germany
- *Andy Georgi,* Technische Universität Dresden, Center for Information Services and High Performance Computing (ZIH), Germany
- *Yenumula B. Reddy,* Department of Computer Science, Grambling State University, USA
- *Wassim Abu Abed,* Institute for Computational Modeling in Civil Engineering, Technische Universität Braunschweig, Germany

INFOCOMP 2013: `http://www.iaria.org/conferences2013/INFOCOMP13.html`

INFOCOMP International Expert Panel: Exa-Intelligence

Related Topics for Audience Discussion:

- Scientific, technical, and high end issues,
- Challenges and experiences from disciplines (natural sciences, technical applications, ...),
- Intelligent and autonomous components,
- Software Defined Networks,
- Parallelisation and localisation,
- High and low level issues,
- Challenges with numerical implementations,
- Suggested application scenarios,
- Knowledge discovery,
- Long-term issues,
- Integration of multi-disciplinary data,
- . . .

INFOCOMP International Expert Panel: Exa-Intelligence

Panel Statements:

- **Key-terms:** Intelligent systems, multi-agent systems, sustainable long-term knowledge discovery, multi-disciplinary context, . . .
- **Intelligent systems:** Very large industrial operations, very large numbers of interacting systems.
- **Learning processes:** Training, ontologies, semantic technologies, on-the-fly, ramp-ups, and learning curves.
- **High End Computing:** Intelligent use of resources.
- **Integrated Information and Computing Systems needed:** Resources for computing and storage (application scenarios in natural sciences, geosciences, archaeology, medicine, . . .).
- **Complexity and decision:** Management processes, stabilisation, (fast!) decision making, modularisation of components and technologies, parallel processes, interfaces, policies.
- **Long-term knowledge resources:** Multi-disciplinary, universal classification, long-term vitality, sustainability, . . .
- **Resources management:** Reduce complexity from planning to operation, with hardware and software.

INFOCOMP International Expert Panel: Exa-Intelligence

Wrapup:

- **Intelligence:** Why intelligent systems?
- **Integration and complexity:** Which challenges do we encounter?
- **Physics and engineering:** What to expect beyond numerics and big data?
- **Computing and storage requirements:** Key issues?
- **Efficiency and resources:** Where to go for efficient modelling?
- **Infrastructures, frameworks, applications:** What do we need?
- **Discovery:** How can ontologies and knowledge discovery tame complexity?
- **Your ideas:** Who and what are we creating and operating high end systems for? Why does general progress take so long?
- **Sustainability:** Perspectives for infrastructures?
- **Call for Collaboration:** Who can contribute and collaborate?

INFOCOMP Expert Panel: Post-Panel-Discussion Summary

Post-Panel-Discussion Summary (2013-11-20):

- Advanced application scenarios drive for more capacity!
- Sustainability with future High End Computing (HEC) resources leads to intelligent application components, e.g., for Agent Based Modelling (ABM), with Exa-operations, Exa-Flop/s, ...
- Future High Performance Computing (HPC), supercomputing, and parallel Multi-Agent Systems (MAS) should become more more outlasting than implementations from scratch for any architecture and system generation.
- Increasing need to provide advanced network support for creating demanding high end services.
- Software Defined Network (SDN) implementation is in development.
- Increased reliability and efficiency regarding operation and emerging technologies is required for Exascale.
- Integrating HEC, MAS, SDN, knowledge resources, and applications.
- Goal for future initiatives and projects is an integrated platform for supercomputing, High End Computing, and mobile computing.

INFOCOMP Expert Panel: Table of Presentations

Panelist Presentations:

- **High End Computing and Advanced Scientific Supercomputing: Sustainability, Challenges, and Prospects with Management and Research** (*Rückemann*)
- **Challenges on Parallelising Multi-agent Systems** (*Leitão*)
- **Exa-Intelligence: On Managerial Challenges in Complex Industrial Landscapes of Risk Case: Production Ramp-ups in Aviation Industry** (*Inden*)
- **Software Defined Networks: Only a hype or the next logical step?** (*Georgi*)
- **Software Defined Networks** (*Reddy*)
- **Exa-Intelligence: A scientific application software's perspective on fault tolerance** (*Abu Abed*)

International Expert Panel INFOCOMP 2013

Exa-Intelligence: Next Generations of Intelligent Multi-Agent and High End Computing Systems in Development and Practice

High End Computing and Advanced Scientific Supercomputing: Sustainability, Challenges, and Prospects with Management and Research

The International Conference on Advanced Communications and Computation (INFOCOMP 2013)
November 18, 2013, Lisbon, Portugal

IARIA

Dr. rer. nat. Claus-Peter Rückemann[1,2,3]

[1] Leibniz Universität Hannover, Hannover, Germany
[2] Westfälische Wilhelms-Universität Münster (WWU), Münster, Germany
[3] North-German Supercomputing Alliance (HLRN), Germany

ruckema(at)uni-muenster.de

Westfälische Wilhelms-Universität Münster

GEXI

ACT MAP

Status: High End Computing and Advanced Scientific Supercomputing

High End Computing . . .

- High End Computing is increaslingy driven by physics (and applications).
- Latency.
- Big data (Volume, Velocity, Variabiltiy, Vitality).
- Long-term storage and archiving.
- Scheduling issues.
- Optimisation and efficiency.
- Interactive high end applications.
- Complementary needs and perception on research and business?
- Integrated collaboration frameworks and concepts for resources, services, and disciplines?
- Validation, verification, error correction?
- Redundancy and criticality management?
- Knowledge, what?, where?
- Classification (Universal Decimal Classification, UDC)?
- Content and context?
- Interactive communication requirements (quantity and quality).
- Data transfer to/from distributed resources (interactive and batch).
- Flexible networks?

Status: High End Computing and Advanced Scientific Supercomputing

Advanced Scientific Supercomputing . . .

- Complexity of computing scenarios is steadily increasing.
- Development of methods and applications depending on funding, physical resources, consulting, reliability, high availability, security.
- Is software a solution for every problem?
- Isolated user groups, no holistic view on context and content.
- Who will be the recipients of YOUR work/results?

Management and Research

Sustainability ...

- Much shorter life-cycle and tender intervals required.
- Much less configuration re-inventing cycles required.
- Implementation of expertise driven management processes required.
- Overall (external) auditing required.

Challenges and Prospects ...

- Measurement is necessary for improvements and discussion.

Challenges of disciplines-services-providers application scenarios

Integrated Information and Computing Systems (GEXI case study)

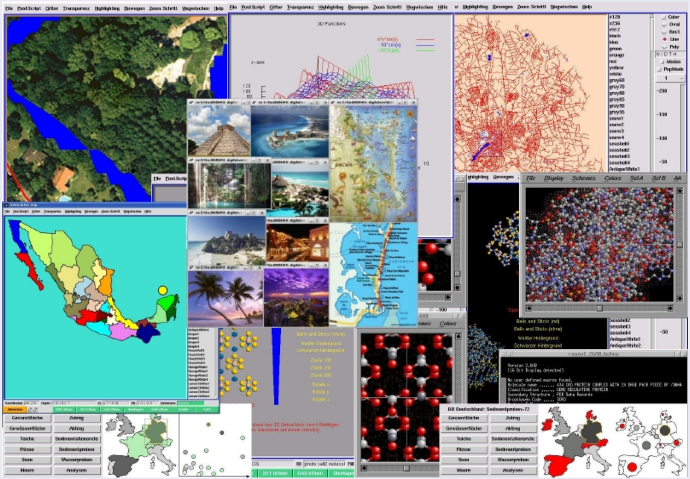

Data objects

- vector
- raster
- aerial
- photo
- spatial
- calculation
- measurement
- processing
- meta objects
- interactive
- commercial
- license
- ...

Universal Decimal Classification

Example excerpt of Universal Decimal Classification (UDC) codes:

UDC Code	Description (English)
UDC 55	Earth Sciences. Geological sciences
UDC 56	Palaeontology
UDC 911.2	Physical geography
UDC 902	Archaeology
UDC 903	Prehistory. Prehistoric remains, artefacts, antiquities
UDC 904	Cultural remains of historical times
UDC 25	Religions of antiquity. Minor cults and religions
UDC 930.85	History of civilization. Cultural history
UDC "63"	Archaeological, prehistoric, protohistoric periods and ages
UDC (7)	North and Central America
UDC (23)	Above sea level. Surface relief. Above ground generally. Mountains
UDC (24)	Below sea level. Underground. Subterranean
UDC =84/=88	Central and South American indigenous languages

Information systems and computed classified objects

Example: Region, Pyramid of Maya, Yucatán, México


Kukulkán


Nohoch Mul


El Meco


El Rey

- **Function:** SAMPLE objects from a group and / or location.
- **Content / context**: compute and storage:
  objects pyramids, Maya, Yucatán region.
- **Computation:** Selection of media photo objects.

Knowledge and computing

Environment and geosciences . . .

- Short-term results favoured.
- Processing.
- Simulation and modelling.

Archaeology and natural sciences . . .

- Documentation and classification.

Vision – Way to go: Intelligent, autonomous components

Intelligent components

- Intelligent software components,
- Knowledge resources,
- Interfaces for least invasive operation,
- Integrated information and computing systems development.

Knowledge resources: (essential transfer over generations)

- Knowledge recognition (expertise) and decision making.
- Knowledge documentation, for any aspect of nature and society (sciences, technical descriptions, tools, cultural heritage, media, ...).
- Long-term means.

Process targets

- Holistic knowledge resources creation,
- Knowledge resources transfer over generations,
- Documentation of requirements and context,
- Integrated information and computing systems development.

Conclusions

Funding (and) multi-disciplinary work and knowledge

Selected challenges and deficits:
(as identified by last years' INFOCOMP Panel on High End Systems)

- Integrating hardware and software solutions!
- Scalability, fast and massive I/O and communication solutions!
- Automation and autonomous components!
- Intelligent components, learning systems!
- Education and teaching!

Search for solutions: Intelligent and High End Computing systems:

- **Multi-disciplinary work** (content, context, knowledge).
- **Sciences** (expertise in different disciplines).
- **Complexity and intelligence** (holistic knowledge and components).
- **No wearout operation with High End Computing and Supercomputing**.
- **Flexible, fast and massive I/O and communication solutions**.
- **Hardware and software** (integration frameworks).
- **Collaboration and funding**.

The Third International Conference on Advanced Communications and Computation

INFOCOMP 2013
November 17 - 22, 2013 - Lisbon, Portugal

Challenges on Parallelising Multi-agent Systems

Paulo Leitão

INSTITUTO POLITÉCNICO DE BRAGANÇA

URL: http://www.ipb.pt
e-mail: pleitao@ipb.pt

Lisbon, 18th November 2013

Contextualization of MAS

- Distributed Artificial Intelligence field
- Society of distributed intelligent agents that interact each other to achieve their goals

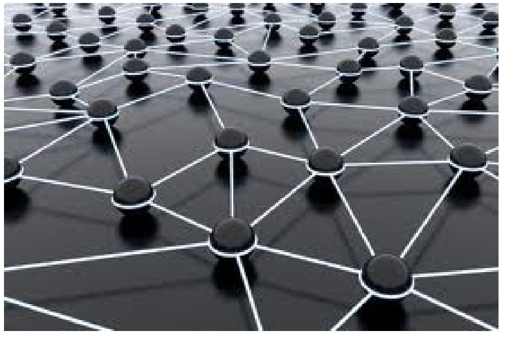

- Suitable to solve large-scale complex problems:
  - Planning and scheduling in manufacturing or logistics
  - Collaborative management of electrical power systems
  - Autonomous traffic systems
  - Computational risk management
  - High level animation and an emergent field in high level simulation
  - ...

Problem

Do we need to run MAS solutions in HPC/HEC platforms?

MAS are distributed systems, but what happen if we parallelise a distributed system?

- Important note:
  - Cluster based distributed memory computing is efficient if the processor cores spend most of the time computing rather than communicating
  - This is however not easy with MAS due to their high levels of communication between agents

Challenges

- Strategies to divide the global problem into smaller ones?
- How to ensure the high level of messages exchanged between distributed agents?
- How to synchronize the parallelized distributed agents?
- How to access huge volume of data in real-time (i.e. without delays to feed the agents during simulation)
- Which strategies to dynamically allocate the needed cores to fulfill the simulation in the given time?
- How to positively influence the direction of simulation making it to converge faster controlling non desired situations on the way?

ABM for Simulation

- Agent-based modelling (ABM):
  - Computational platform to analyse, experiment and simulate systems populated by agents
  - Reproduction of complex phenomena, such as evolution and self-organisation
  - Examples: Netlogo, Mason and Repast

- Differences to MAS frameworks:
  - MAS frameworks allow building agent-based systems, but they haven’t a simulation infrastructure (e.g., misses a scheduler and the notion of a "clock")
  - ABM frameworks allow agent-based simulation but they have not the purpose of developing agent systems (nor FIPA compliant)

REPAST HPC

- Written in C++ and MPI for parallel operations through distributed memory computing
- Features for the parallelisation of the agent-based model
  - Synchronisation scheduling of events
  - Global data collection
  - Automatic management of agent interactions across multiple processes
- Supporting cross-process communication and synchronisation of the simulation
  - distributed across multiple processes (each containing several agents)

ABM HPC or parallelisation of MAS from scratch?

Other Questions

- Are MAS based simulation technics valuable in a distributed global system, where normal endpoints (such as our laptops) are the cores of an HPC?

  [see for instance the SETI@home project]

- What will be the usage of mobile devices in the future cloud HPC computing?
  - Can this be considered as a future trend?
  - Will this enable the ubiquitous computing and enable a greater level of MAS based cloud computing?

Exa-Intelligence:

On Managerial Challenges in Complex Industrial Landscapes of Risk

Case:

Production Ramp-ups

in Aviation Industry

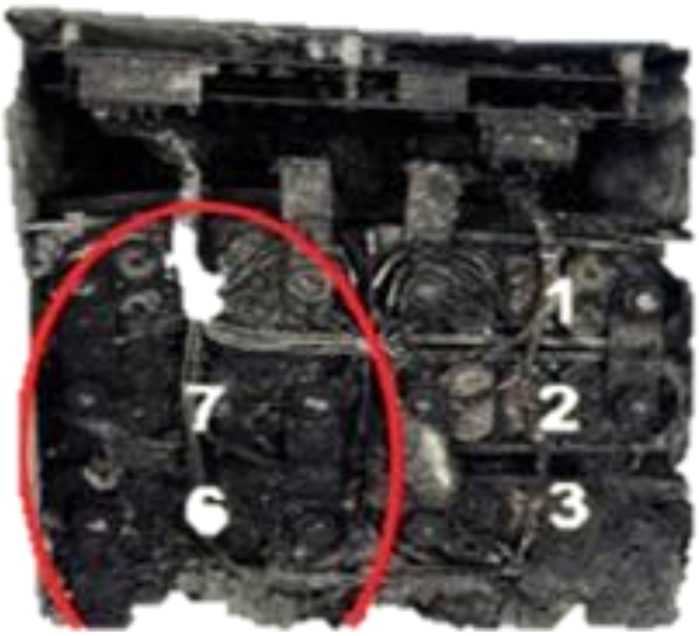

On **Computational** Challenges in Complex Industrial Landscapes of Risk

Some Theses

1. **Demand for computational power in real** (and digital) **economy is crucial to drive and keep step with change.**
2. **Capacity is "going service"**
   In far most of applications not "power" will be the problem but architectures making capacity available for solutions.
   (Sources e.g.: Panels Infocomp 2011 / 2012, FP/8 Horizon 2020)
3. **The challenge is to adjust architectures of capacity demand and of capacity services to each other – e.g. in terms of**
   - **Computation / communication in processing**
   - **Homogeneity / heterogeneity of the problem**

Heterogeneous
End-to-end operations' networks
(large scale of different activity/objects)

Homogeneous
Continental weather model
(multiple identical cells)

Making homogeneous
Translating Operations networks into a landscape of operations' risk

Distributed Operations' and Operations' Management System

Landscape of Coupled Risks across Factories (the whole picture)

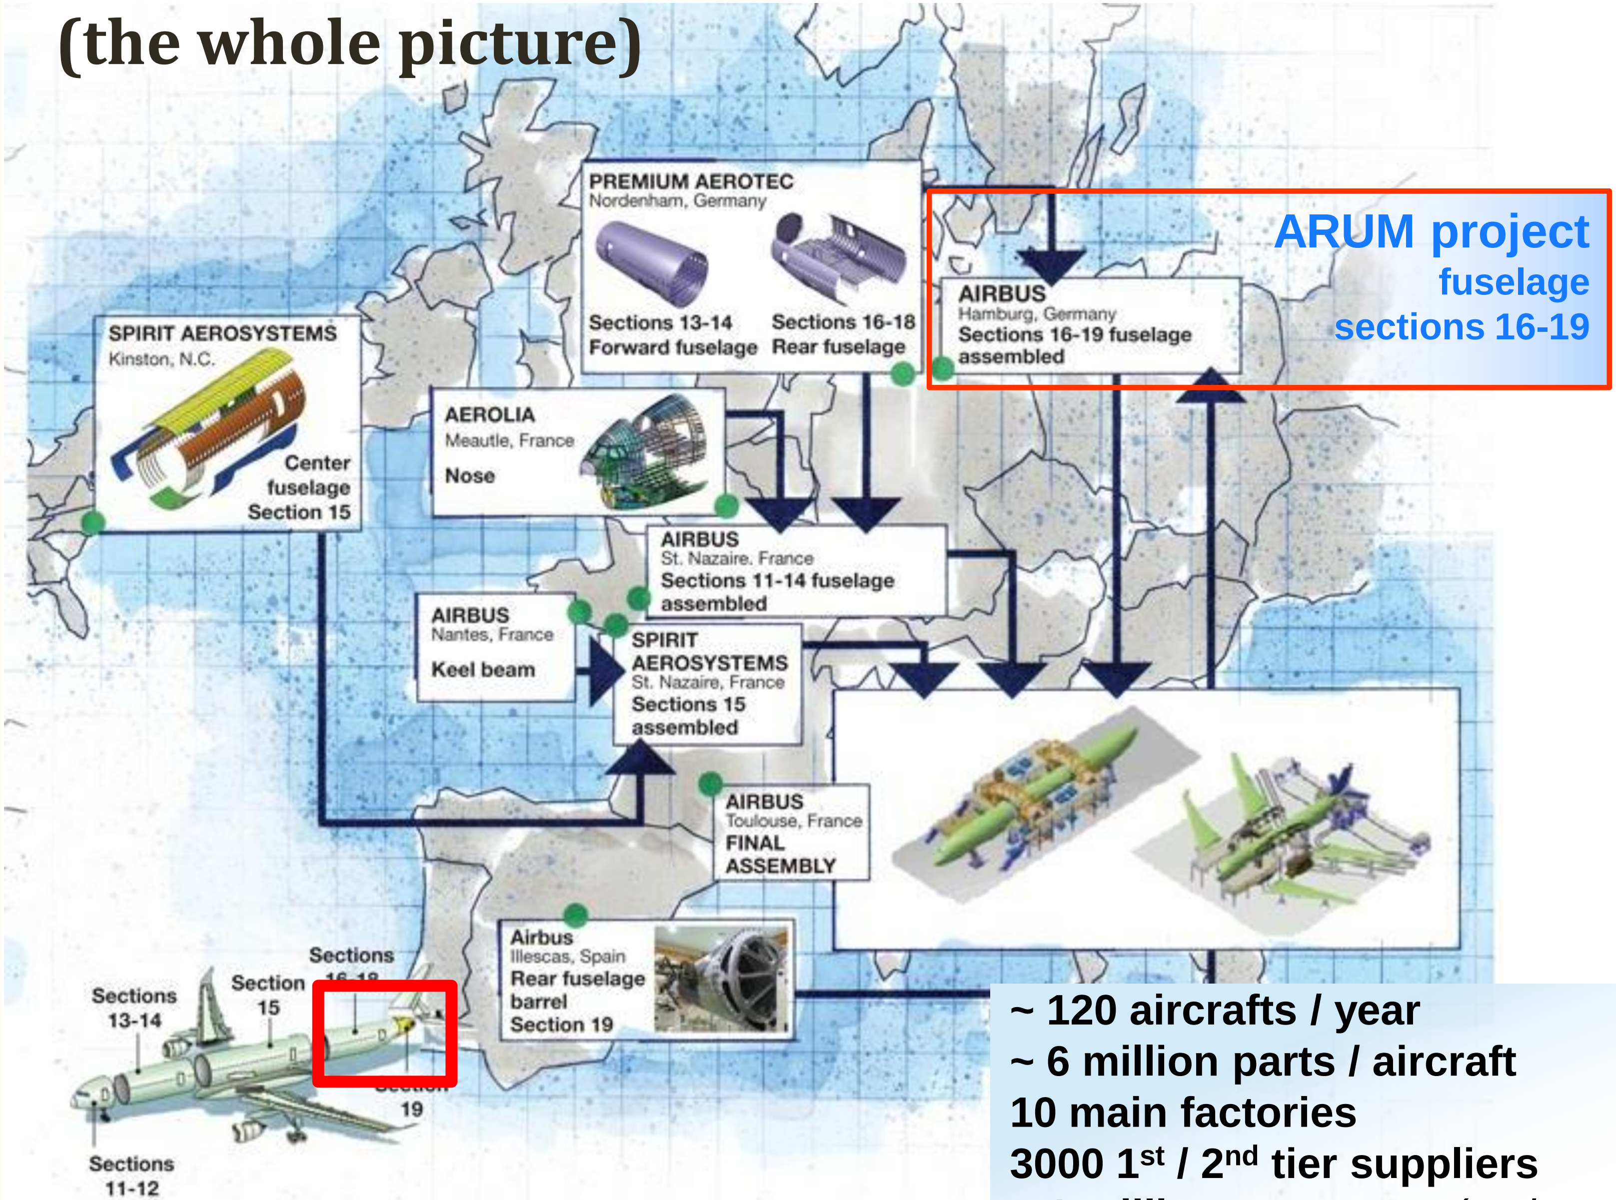

~ 120 aircrafts / year
~ 6 million parts / aircraft
10 main factories
3000 1st / 2nd tier suppliers
> 1 million processes (*estimate*)

http://d2n4wb9orp1vta.cloudfront.net/resources/images/cdn/cms/0911hpc_Airbus_Fuselage_drawing.jpg

Udo Inden, Cologne University of Applied Sciences

Distributed Operations' and Operations' Management System

Landscape of Coupled Risks in a Factory

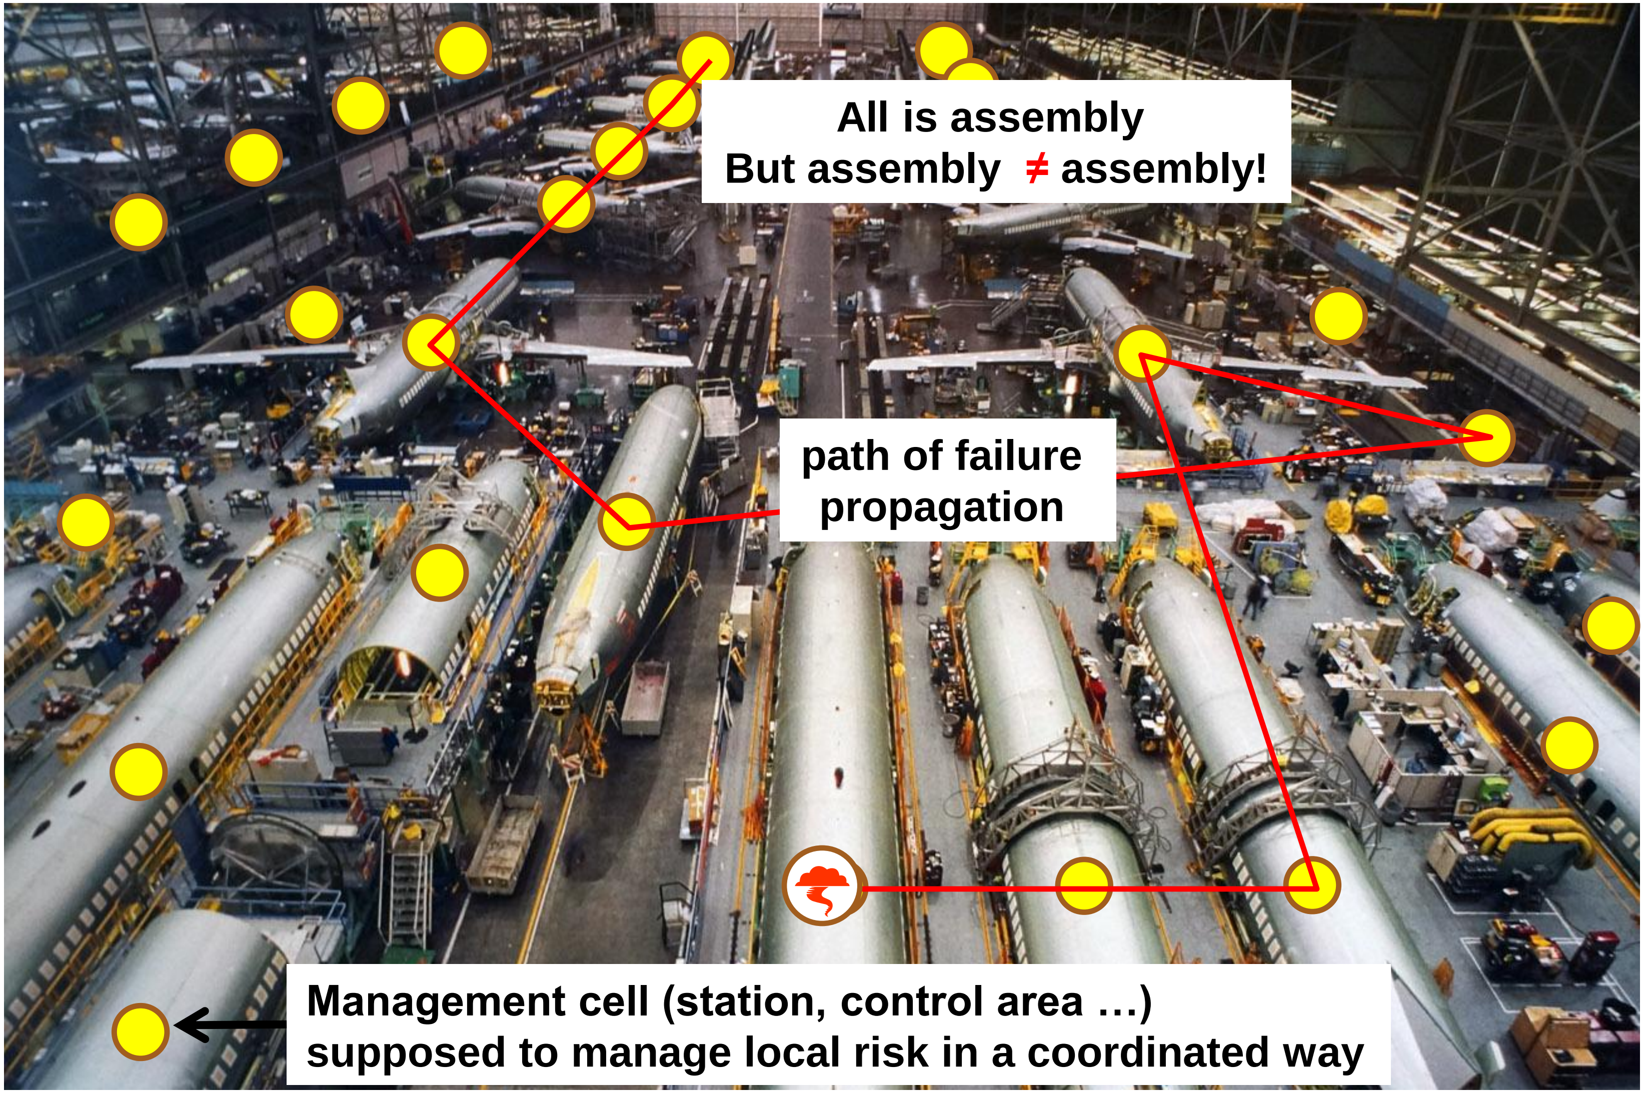

Boeing 737 assembly line in the 1980s and 1990s, Renton, WA

Core Activity of Operations' Management

R.E.A.L. → Realise – Evaluate – Act – Learn

Realise unplanned event

monitor operations, anticipate, identify, categorize

Evaluate impact of event and calculate relevance

estimate expectation value of propagation, compare with threshold of relevance

$$EV_{\alpha P} = \sum_{\alpha L} r_{\alpha L} * I_{\alpha L}$$ decision criterion: $EV_P > T_R$

P = propagation , **EV**$_P$ = expectation value. **α** = event, **r** = event risk, **I** = impact, **L** = location
R = relevance, **T**$_R$ = threshold value of relevance

Act on relevant events

identify / analyse / plan / schedule / implement options, control effectiveness

Learn how to stabilise system behaviour

find, structure and tune relevant **control parameters** to achieve regular behaviour

The Model Underlying R.E.A.L.

Interdependent Multi-dimensional Lattice Trees

Link = estimated event risk

Node = "management cell" = area of responsibility

Size = estimated impact (effect at aiming point)

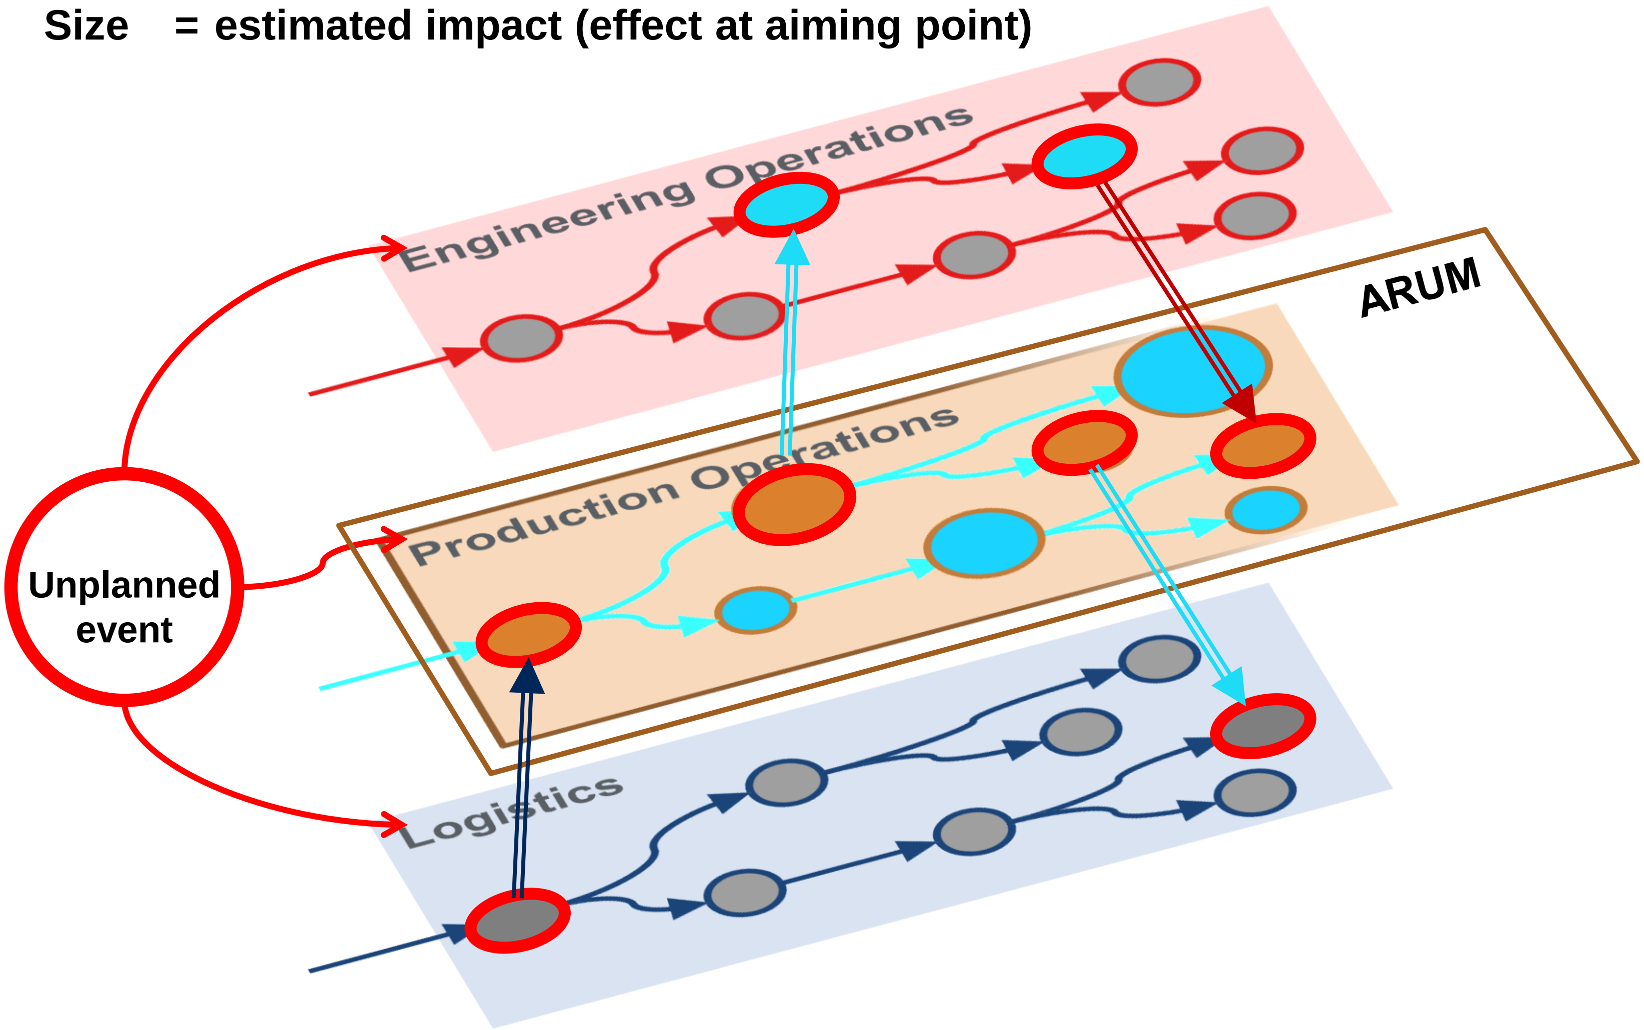

TECHNISCHE UNIVERSITÄT DRESDEN

Zentrum für Informationsdienste und Hochleistungsrechnen – TU Dresden

Software Defined Networks
Only a hype or the next logical step?

Exa-Intelligence: Next Generations of Intelligent Multi-Agent and High End Computing Systems in Development and Practice

Andy Georgi

18. November 2013

Nöthnitzer Straße 46
01187 Dresden

Telefon: +49 351 - 463 38783
E-Mail: Andy.Georgi@tu-dresden.de

Center for Information Services & High Performance Computing

Networking Today

Infrastructure characteristics:

- Static configuration
- Best-effort service
- Over-provisioning
- Single domain management

Traffic characteristics:

- Highly individual requirements
- Unpredictable dynamic flows
- Low utilization (on av.)
- Traffic across multiple domains

Divergence between infrastructure and traffic characteristics!

Software Defined Networking

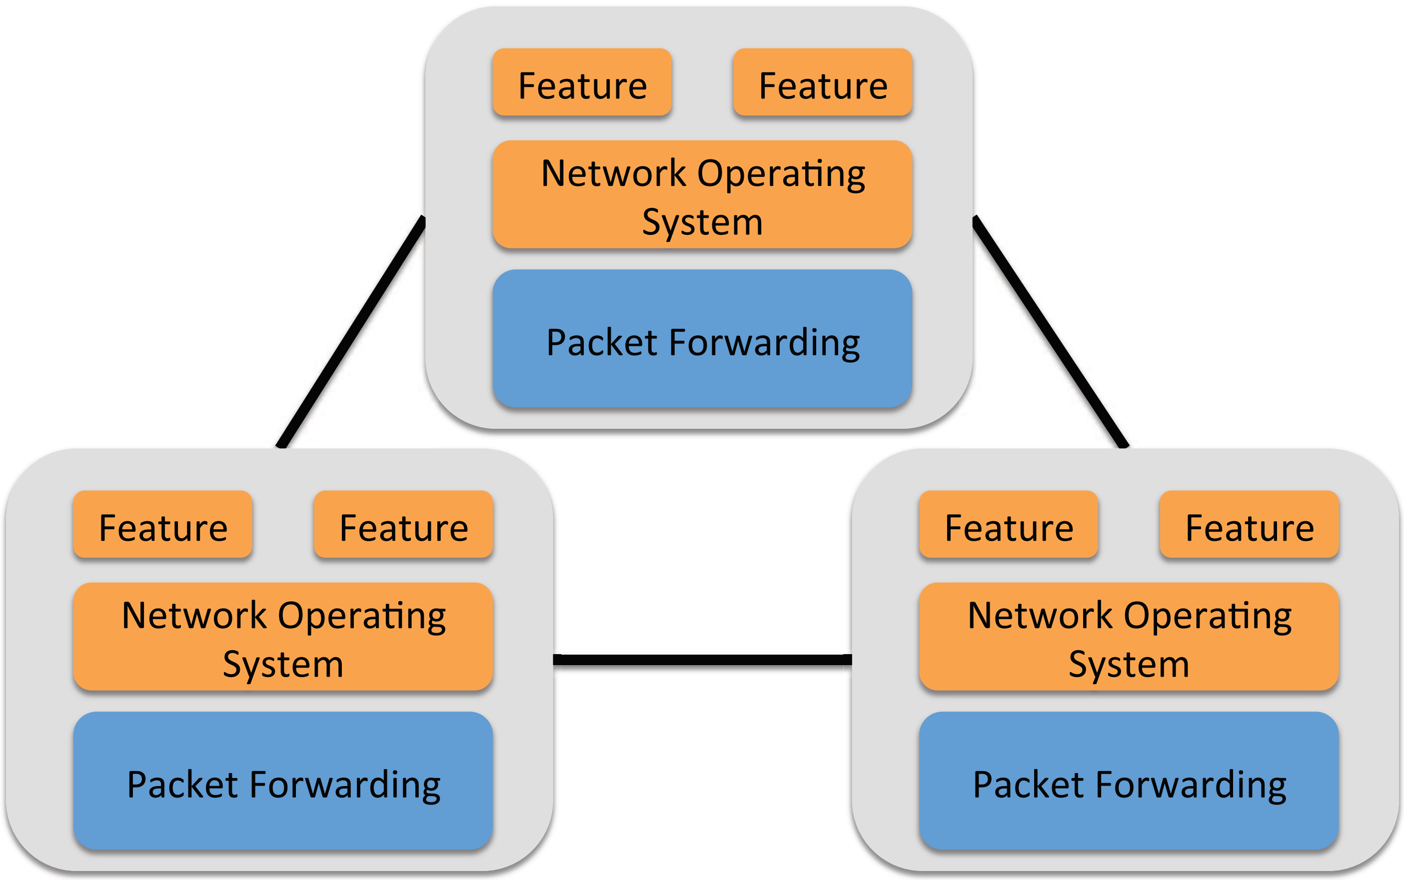

Software Defined Networking

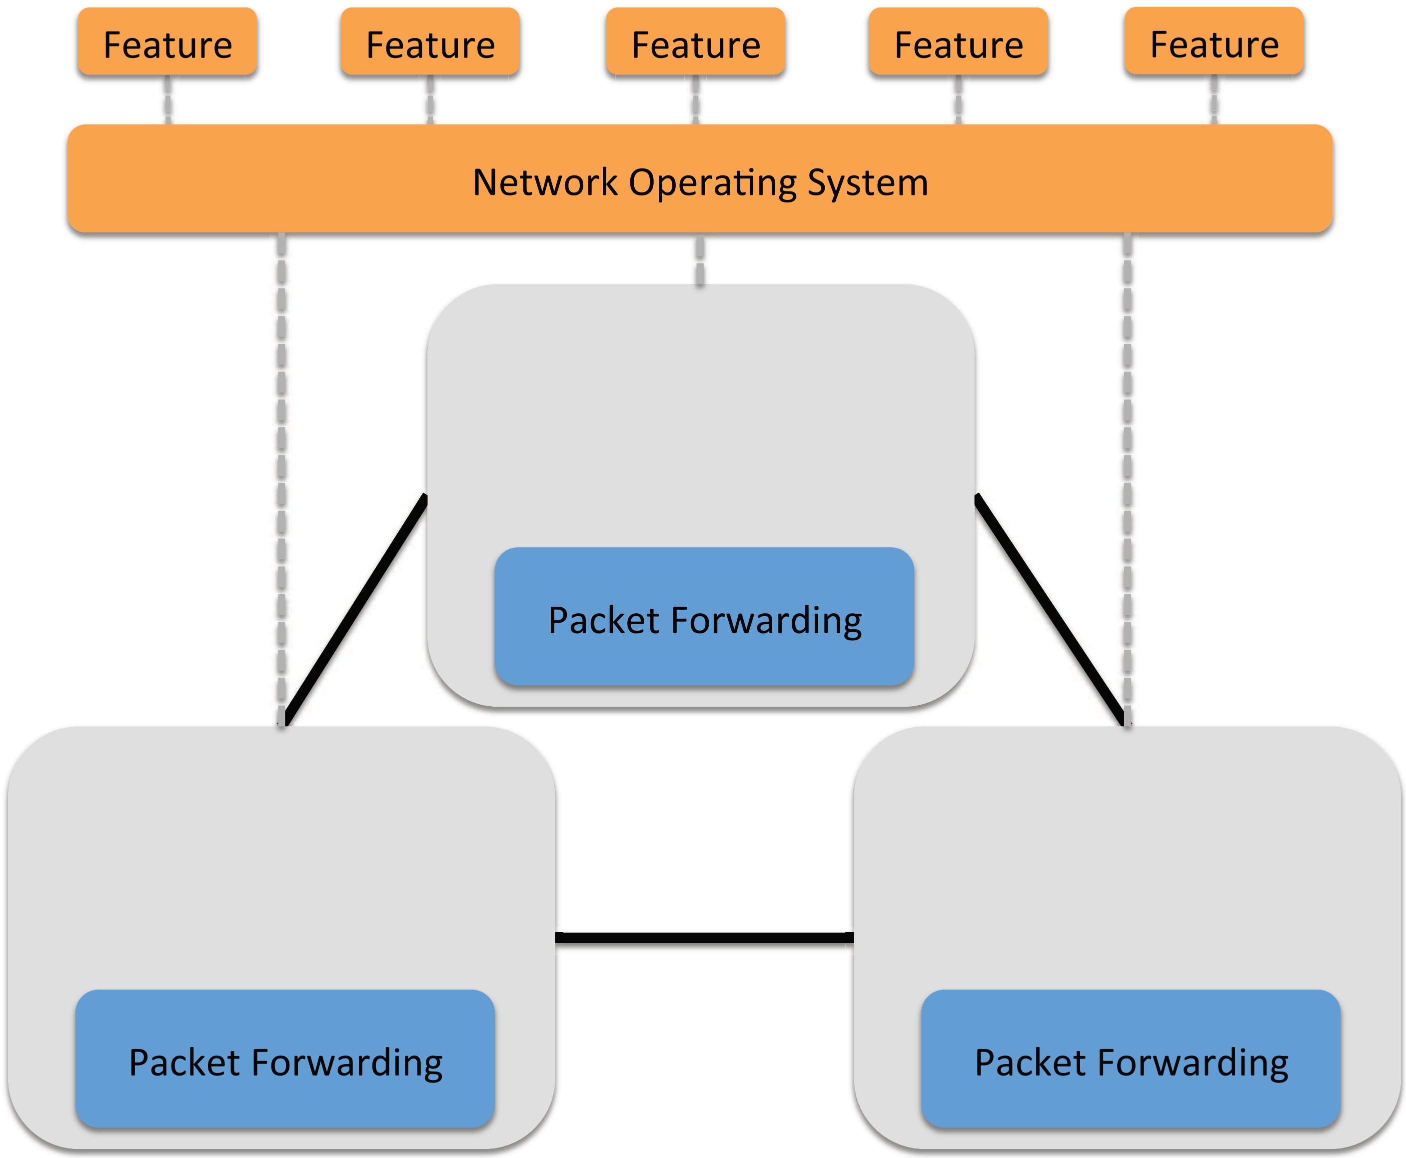

TECHNISCHE UNIVERSITÄT DRESDEN

ZIH Center for Information Services & High Performance Computing

Future of Networking

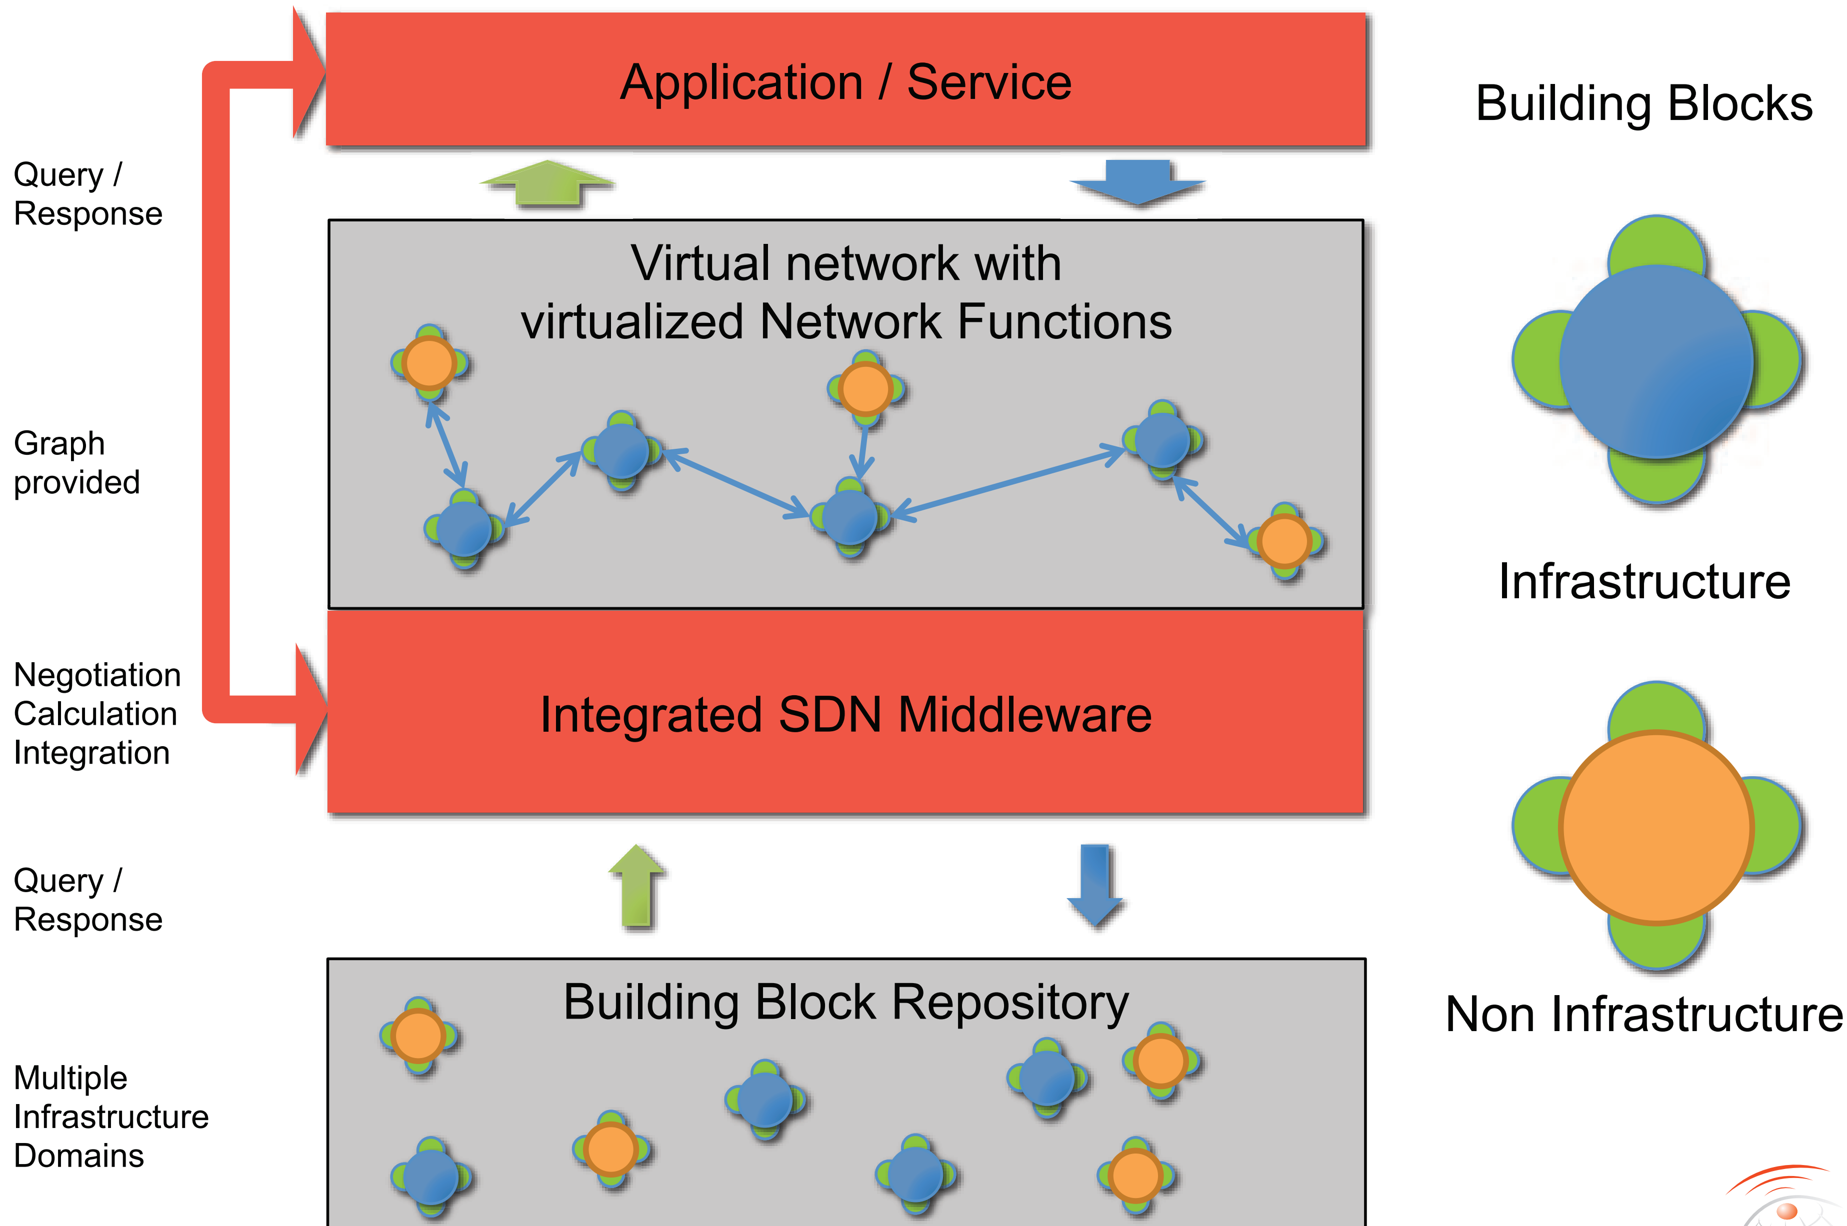

TECHNISCHE UNIVERSITÄT DRESDEN

ZIH
Center for Information Services & High Performance Computing

Thank you for your attention!

Software Defined Networks

Yenumula B Reddy

Grambling State University

Need of Software Defined networks
Limitations of Current Network

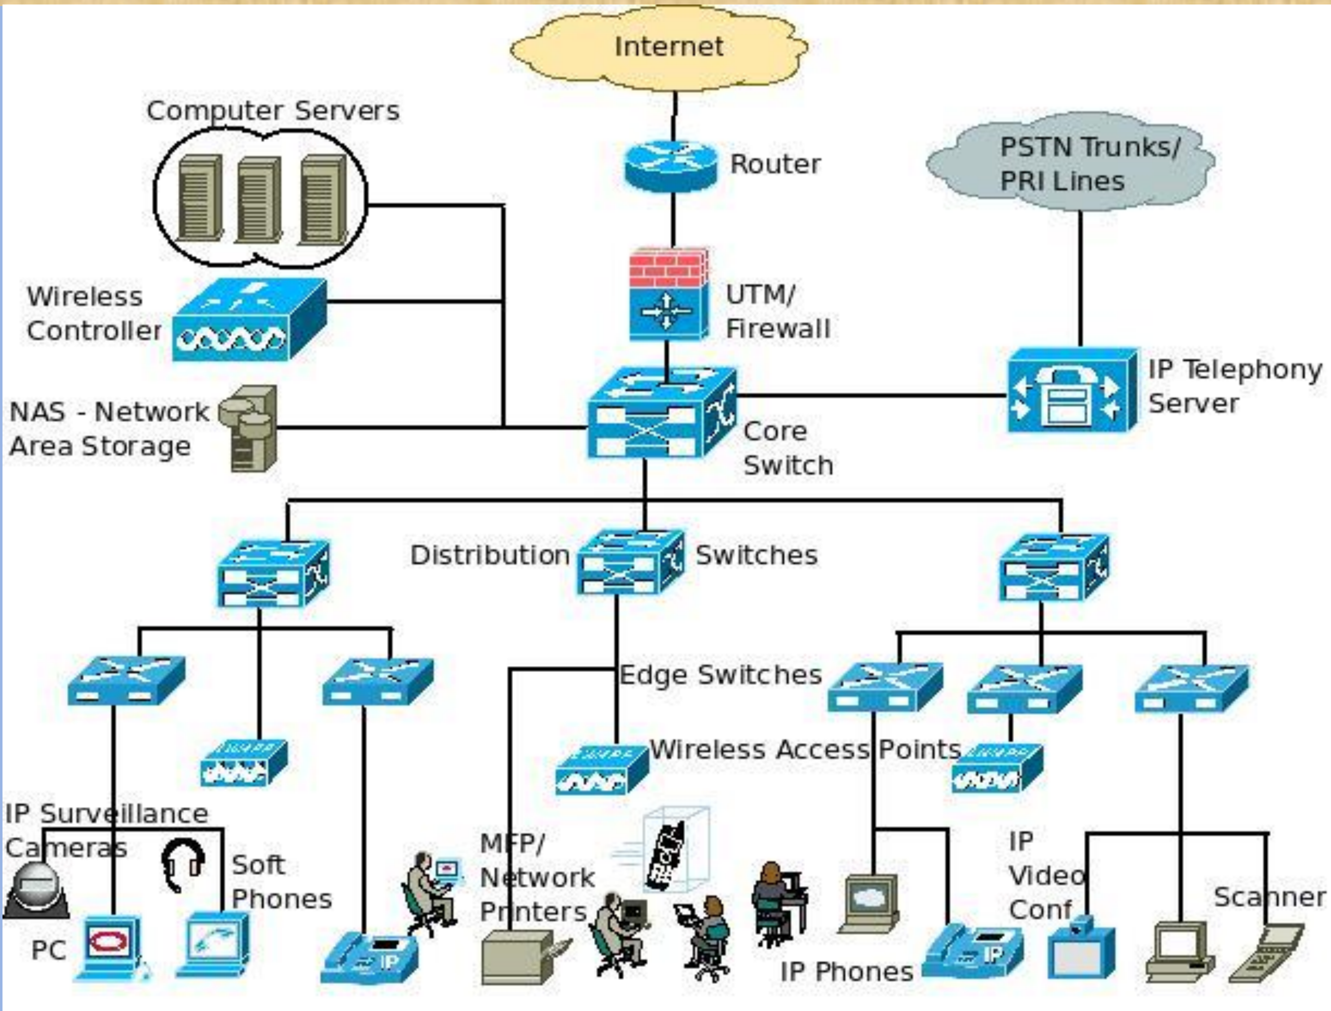

Major components connected in the network

Internet - Internet cloud refers to the source of the Internet to an organization

Router - as a gateway between the LAN and the WAN networks

Unified Threat Management (UTM)/ Firewall - for providing gateway level network security for the various end points used in the organization

.**Core Switch – a layer 3 based network switch** - a Layer-3 based Network Switch that connects to the various distribution switches, edge switches using Optical Fiber Networks or UTP Copper cabling.

Network Area Storage (NAS) Device – to store bulk data

Wireless Controller - to provide wireless (Wi-Fi) access to the PC's/ Laptops etc

.**IP Telephony Server** - Servers provide centralized administration and connectivity to PSTN Lines to all the IP Phones/ VOIP devices

Distribution Switches - provide an aggregation layer for network switching.

Edge Switches - are basically Layer-2 switches that provide direct connectivity to the various network devices like PC's, laptops, Wireless Access Points etc using the Copper UTP cables.

Wireless Access Points - contain built-in radios which provide wireless signals for connecting certain network devices that has an in-built wireless adapter.

Network Endpoints/ Devices – includes various network devices/ endpoints connecting to the LAN via edge switches/ wireless access points.

http://www.excitingip.net/27/a-basic-enterprise-lan-network-architecture-block-diagram-and-components

Limitations _Contd

- **Difficult to manage the current network traffic and increasing demands**
- **Can not change dynamically the configuration according to network conditions (Reason: many complex functions backed into infrastructure)**
- **No control plane abstraction for the whole network**
- **No network operating systems to have global view of network**

Need of Software Defined networks

Software defined networks (SDN)

Definition:

- Software-defined networking (SDN) is an approach to networking in which control is decoupled from hardware and given to a software application called a controller.

- The goal is to respond quickly to meet the business requirements
- Can maintain the traffic from centralized control without having to touch the switches.
- Can change the rules, prioritize the blocks to manage the traffic

Need of Software Defined networks

Overcoming the problems with SDN

- **Separate control plane and Data plane entities**
- **Run control plane software on general purpose hardware**
- **Programmable data planes**
- **An architecture to control entire network**

Need of Software Defined networks

Benefits of software defined network

- **Lower operational expenses**
- **Flexibility – service providers can develop their services by using standard tools**
- **Improved uptime by eliminating manual intervention**
- **Better management for service provides through virtual network**
- **Planning the resources through better visibility into network**
- **Infrastructure savings by separating control pane and data plane**

Need of Software Defined networks

Promises of software defined network

- **Emerging network architecture to unlock the innovation and upgrade the networking efficiencies**
- **It is the network that decoupling of the network control layer from the data transport layer.**
- **It is dynamic, adaptable, manageable and cost-effective**
- **Cloud computing is also driving the need for SDN**
- **Facilitate network automation, virtualization, and policy management**

Need of Software Defined networks

Current status

- The IETF is investigating models of SDN for feasible technical approaches
- Simultaneous to SDNRG's study, other IETF (internet emerging task force) Working Groups have started their own efforts in SDN
- SDN, especially OpenFlow, has already been used in a carrier's production network to provide virtualized ne
- Until recently, the IP/MPLS network settings of tunnel LSPs to carry user's traffic were statically and manually configured at routers. As a result, it was technically difficult to deal with a scenario in which source-destination pairs of LSPs frequently changed. SDN's logically centralized approach has the capability to solve this kind of technical challenge.

Challenges

- **Will it change the nature of IT**
- **Does it requires resources around cloud?**
- **Do we need better support of cloud deployments**
- **Can it Control entire data from single location**
- **Will SDN help to spin up new applications**
- **Will it enhance the IT security**
- **What are the design challenges**
- **Does it Require more intelligence in decisions**

Future of Software Defined Networks

- SDN is still in its infancy
- Evolutionary approach to network architecture - programming network devices to modify their behavior
- Amazon and Google are already starting to use SDN
- Future development of new applications and services that make the most of SDN
- The development of new applications and services that make the most of SDN
- Impact on hardware manufacturers and Business
- High security level at networking – user control

Technische Universität Braunschweig

iRMB

Exa-Intelligence: A scientific application software's perspective on fault tolerance.

Wassim Abu Abed, PANEL INFOCOMP 2013, Nov. 18. 2013, Lisbon

Status: Fault-tolerance (FT) on peta-scale

Almost all FT Techniques are based on a global checkpoint-restart recovery model

Indiscrimination between different types of faults

- Permanent node crash
- Detected transient errors
- Network errors
- File system failures

Some exceptions exist, but:

- Sporadic attempts in some components
- Higher level of software stack does not cope with faults (Runtime system)
- No fault detection and management across the software stack
  - Due to lack of communication and coordination between
    - software layers and components
    - Application and supporting libraries

Status: Fault-tolerance on exa-scale

Checkpoint-restart recovery model
on exascale systems
WILL NOT BE ENOUGH !

Why:

- Increased failures.
- Time for CP/R will exceed the mean time to failure of system.
- Possibly no application will run to the end.
- Increased runtime/power costs due to CP/R. (CP/R is expensive)
- Hardware level reliability improvement is expensive.

Vision: Intelligent exascale fault tolerance

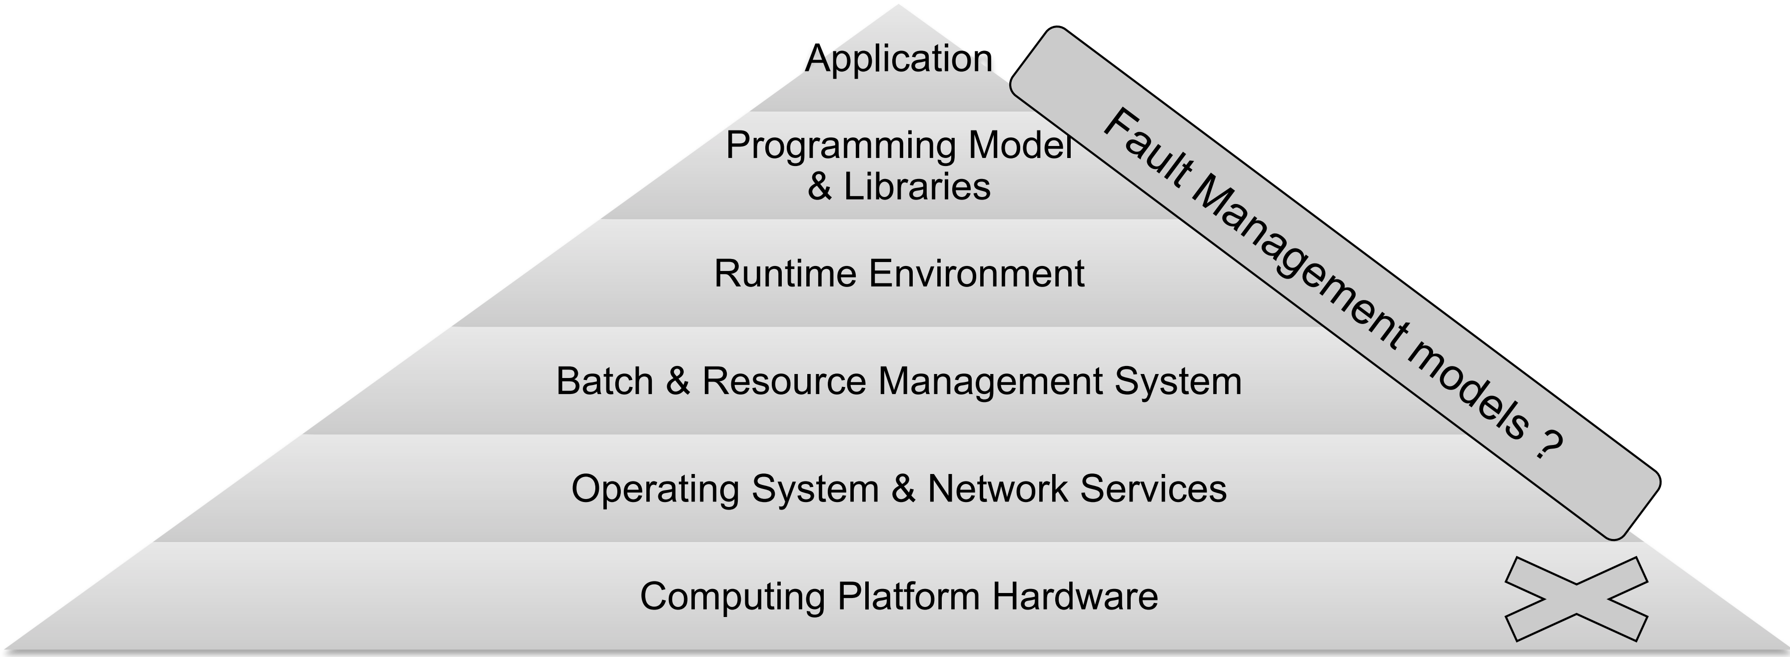

Intelligent Fault Tolerance Mechanism

Goal: Running applications until their normal termination, despite the essentially unstable nature of exascale systems.

References:

- M. A. Heroux, Toward Resilient Algorithms and Applications, Invited Talk on FTXS 2013, New York City, June 18th 2013.
- F. Cappello, A. Geist, B. Gropp, L. Kale, B. Kramer, and M. Snir. Toward exascale resilience. The International Journal of High Performance Computing Applications, 23(4): 374–388, 2009.
- E. Dubrova. Fault-Tolerant Design. Springer, 2013.
- J. Daly (eds.), Inter-Agency Workshop on HPC Reilience at Extreme Scale, National Security Agency Advanced Computing Systems, Februrary 2012.
- W. Bland, A. Bouteiller, T. Herault, J. Hursey, G. Bosilca and J. Dongarra, An Evaluation of User-Level Failure Mitigation Support in MPI, Proceedings of Recent Advances in Message Passing Interface 19th European MPI Users Group Meeting, EuroMPI 2012. Springer, Vienna, Austria, Sep. 2012, pp. 193–203.
- W. Bland, P. Du, A. Bouteiller, T. Herault, G. Bosilca and J. Dongarra, A Checkpoint-on-Failure Protocol for Algorithm-Based Recovery in Standard MPI, Kaklamanis et al. (eds.) Euro-Par 2012, LNCS, vol. 7484, Springer-Verlag, Berlin Heidelberg, 2012, pp. 477–488.
- Z. Chen and J. Dongarra, Algorithm-Based Fault Tolerance for Fail-Stop Failures, IEEE Transactions on Parallel And Distributed Systems, Vol. 19, No. 12, 2008, pp. 1628–1641.