From

Infinite Do Loops

to

Never Work Again

On the occasion of the retirement of

Henri Bal

by

Cees de Laat

Some first remarks

- This is not a scientific talk!
- But some serious and funny remarks

ICT to support the transformation of Science in the Roaring Twenties

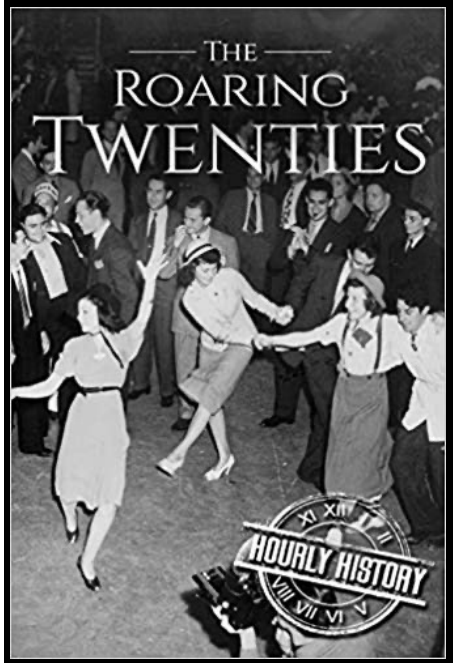

From Wikipedia: The Roaring Twenties refers to the decade of the 1920s in Western society and Western culture. It was a period of economic prosperity with a distinctive cultural edge in the United States and Western Europe, particularly in major cities such as Berlin, Chicago, London, Los Angeles, New York City, Paris, and Sydney. In France, the decade was known as the "années folles" ('crazy years'), emphasizing the era's social, artistic and cultural dynamism. Jazz blossomed, the flapper redefined the modern look for British and American women, and Art Deco peaked….

This period saw the large-scale development and use of automobiles, telephones, movies, radio, and electrical appliances being installed in the lives of thousands of Westerners. Aviation soon became a business. Nations saw rapid industrial and economic growth, accelerated consumer demand, and introduced significantly new changes in lifestyle and culture. The media focused on celebrities, especially sports heroes and movie stars, as cities rooted for their home teams and filled the new palatial cinemas and gigantic sports stadiums. In most major democratic states, women won the right to vote. The right to vote made a huge impact on society.

Some jokes

- Kings and Bishops
- From this Day

Earliest Henri encounters 1996

- **DAS-1**
- **ATM**
- **We did not have a cluster in Utrecht**
- **Not yet part of ASCI**

To Boston - Digital Equipment

Date: Mon, 17 Jun 1996 20:35:57 +0200
To: peterslo@fwi.uva.nl (P.M.A. Sloot), rvd@nikhef.nl, miron@cs.wisc.edu, bjo@fwi.uva.nl, joep@fwi.uva.nl, frank@fwi.uva.nl, berry@fwi.uva.nl, epema@twi.tudelft.nl, ar@uni-paderborn.de, emmen@genias.nl, W.Lourens@fys.ruu.nl, C.T.A.M.deLaat@fys.ruu.nl, pier@nikhef.nl, wielinga@sara.nl, bal@cs.vu.nl, ast@cs.vu.nl
From: C.T.A.M.deLaat@fys.ruu.nl (C. de Laat)
Subject: Re: Meeting June 21th

➢ I herewith suggest the following agenda for the upcoming Distributed MetaComputing meeting: ….

Since I will be in the USA from Wednesday 19th until Sunday 23th on a business trip I will be unable to attend this meeting.

regards,

Cees de Laat.
Faculty of Physics and Astronomy, Utrecht University,
Princetonplein 5, NL-3584CC Utrecht, The Netherlands.
delaat@fys.ruu.nl , Tel: (31)30-2534585

DAS developments

- DAS-1 Wide-area computing (1997)
  - Features: Homogeneous hardware and software, dedicated ATM network
  - 6 Mb/s per link
  - Nowadays with that bw your children would divorce from you!
  - In 1995 the entire university had 4 Mb/s, later 34 Mb/s and 155 Mb/s
    - 1997 report: We tested this ATM network and found that we always achieve the throughput of about **5 MBit/s** between nodes with constant round trip times of about **3 msec**. We compared this to normal ethernet-router-SURFnet connections and found that the throughput averaged to about **1.5 MBit/s** and round trip times of about **10 msec**. We therefore concluded that the ATM connections deliver a predictable service which is about **ten times better** for our purposes **than** the **normal** university connections.
  - In 1999 stopped ATM research, switched to Ethernet.
  - Standard interface to universities
  - Sonet/SDH in WAN to StarLight
  - Protocol research
  - Lambda Workshop 9/11and iGrid2002

DAS developments

- DAS-1 Wide-area computing (1997)
  - Features: Homogeneous hardware and software, dedicated ATM network
- DAS-2 Grid computing (2002)
  - Features: Globus middleware
  - Internally Myrinet and Fast Ethernet
    - (Paola and I visited Myrinet in Pasadena in January 2006)
  - Gigabit Ethernet to WAN
  - Lots of mails about the extra slot it would cost in head nodes ☺

DAS developments

- DAS-1 Wide-area computing (1997)
  - Features: Homogeneous hardware and software, dedicated ATM network
- DAS-2 Grid computing (2002)
  - Features: Globus middleware
- DAS-3 Optical Grids (2006)
  - Features: Photonically switched 10 Gb/s links between all sites

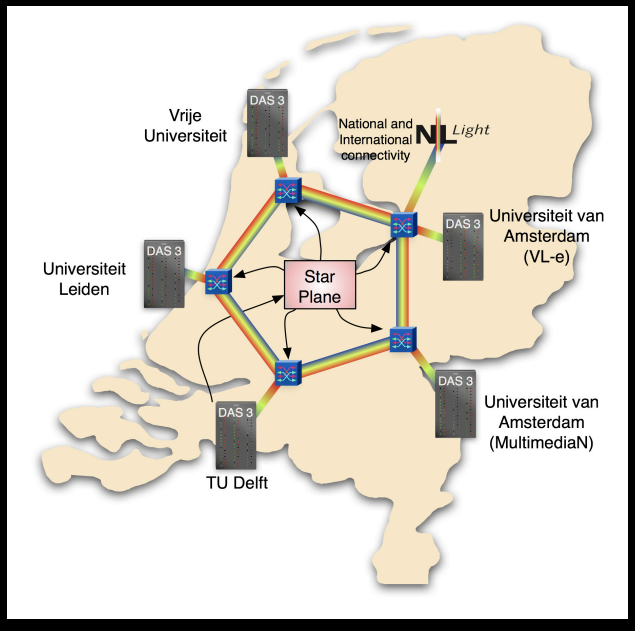

Pict by Paola Grosso

Amsterdam CineGrid S/F node "COCE"

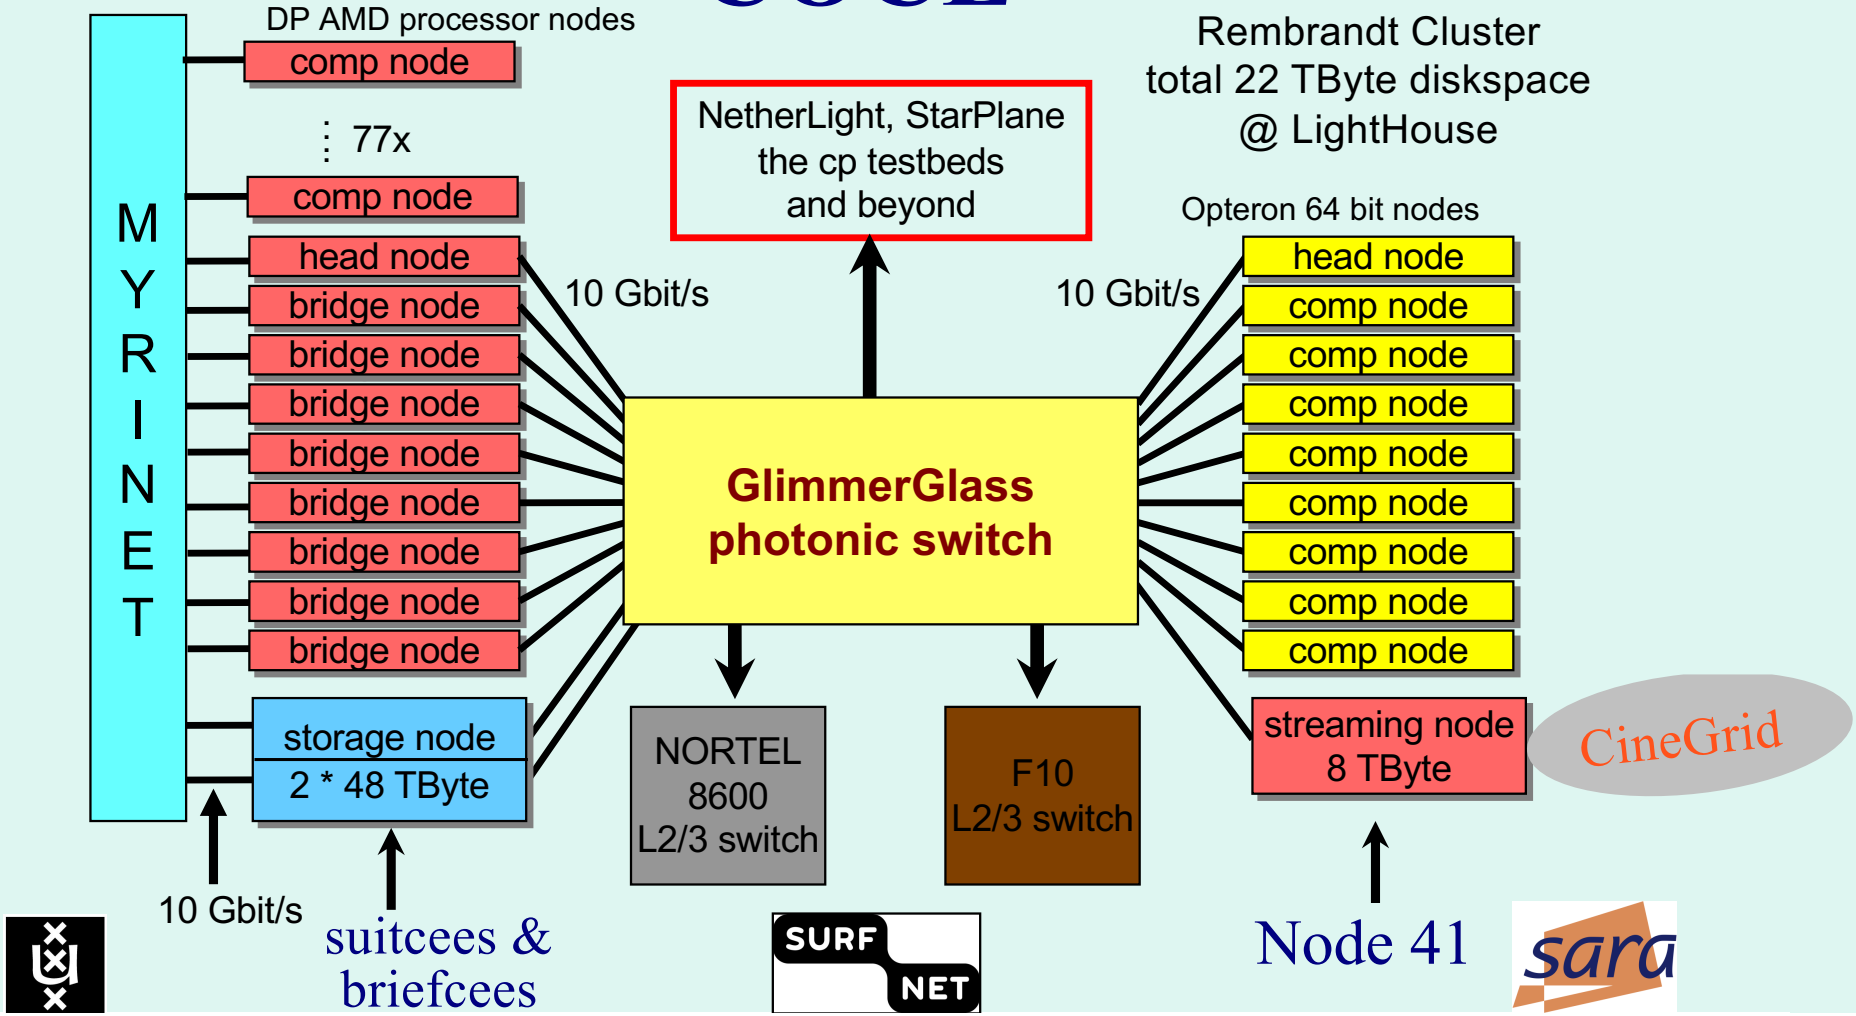

Traveling with Cees (HPDC'2010, Chicago)

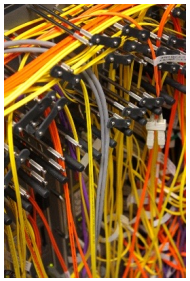

Networking with Cees

From speech by Henri on my retirement

My recollection of that trip

- Joe Mambretti

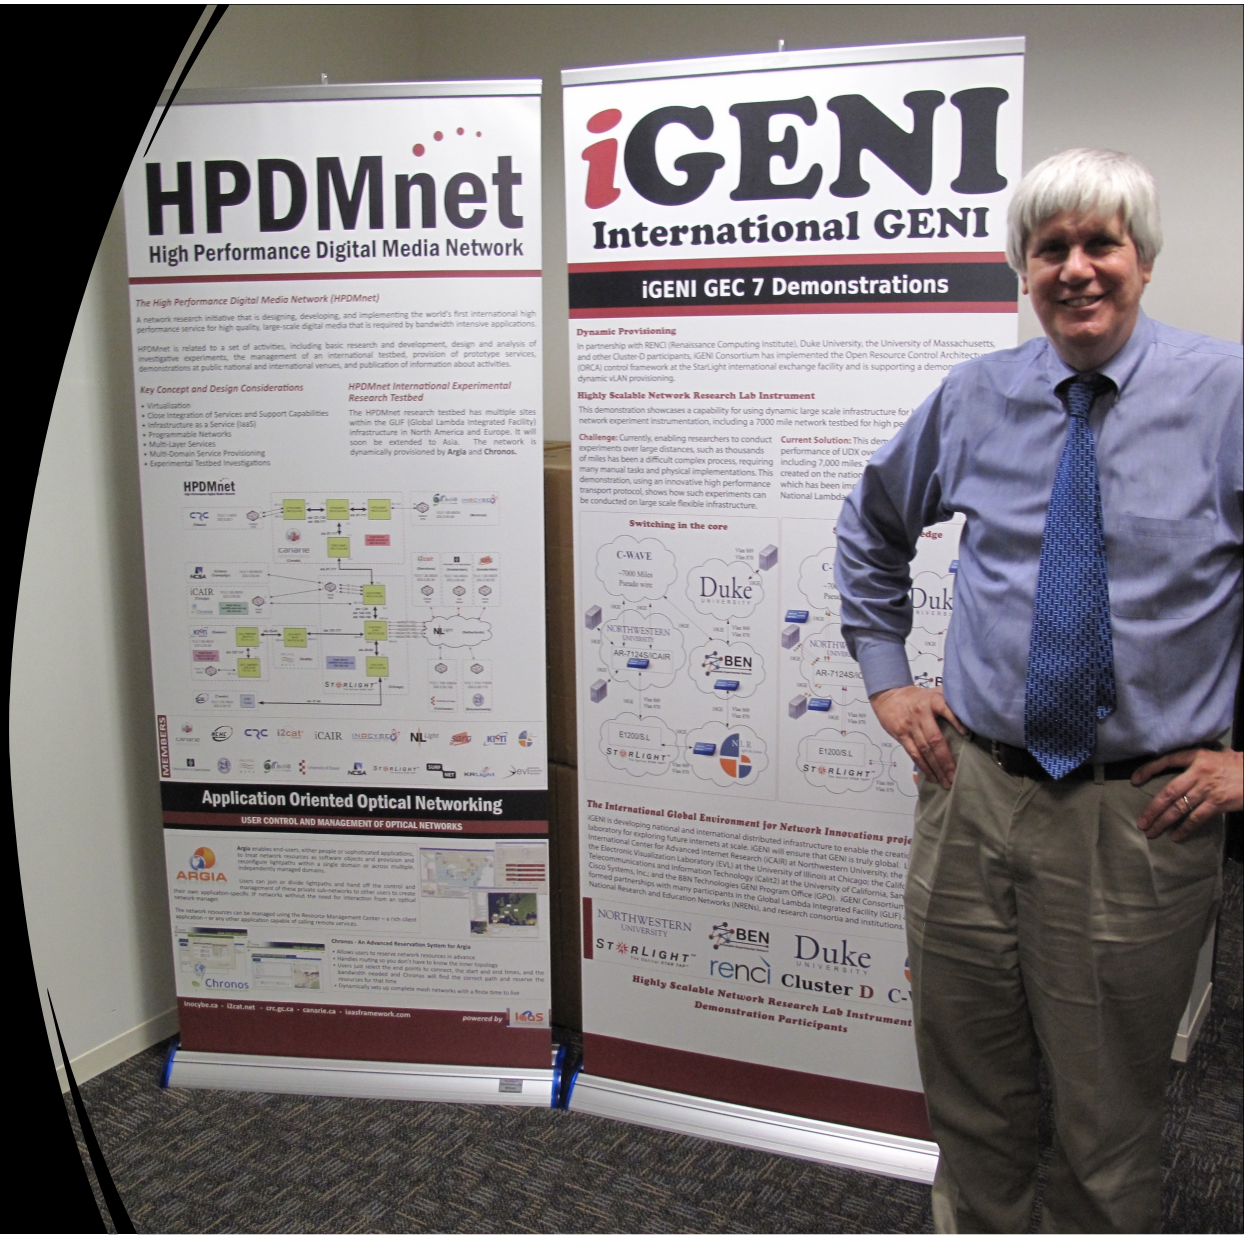

My recollection of that trip

- Joe Mambretti
- His office (Lake Michigan)

My recollection of that trip

- Joe Mambretti
- His office (Lake Michigan)
- StarLight

My recollection of that trip

- Joe Mambretti
- His office (Lake Michigan)
- StarLight
- Surprising ☺

DAS developments

- DAS-1 Wide-area computing (1997)
  - Features: Homogeneous hardware and software, dedicated ATM network
- DAS-2 Grid computing (2002)
  - Features: Globus middleware
- DAS-3 Optical Grids (2006)
  - Features: Photonically switched 10 Gb/s links between all sites
- DAS-4 Clouds, diversity, green IT (2010)
  - Features: Hardware virtualization, accelerators, energy measurements
- DAS-5 Harnessing diversity, data-explosion (2015)
  - Features: Wide variety of accelerators, larger memories and disks, software defined networking
- DAS-6 Distributed trusted data analytics (2019)
  - Features: data stream in network processing, accelerators, virtualisation, security and trust

DAS from Homogeneity to Diversity

	Year	PhDs	Research agenda	DAS Infrastructure
DAS-1	1997	7	Wide-area computing	4 fully homogeneous clusters with ATM network
DAS-2	2002	22	Grid computing	5 homogeneous clusters, Globus, Internet
DAS-3	2006	36	Optical grids	5 clusters with dedicated light paths
DAS-4	2010	33	Clouds, diversity, green IT	6 clusters, virtualization, accelerators, energy
DAS-5	2015	40	Harnessing diversity & complexity	6 clusters, variety of many-core accelerators
DAS-6	2019	>65	Distributed trusted data analytics	6 clusters with different functionalities

Some progress (SC18 exhibit)

2018

? 540 MHz
? MFlops
1000 MByte memory
16000 MByte ssd
0,0012 kWh – 18 h

= ~7×

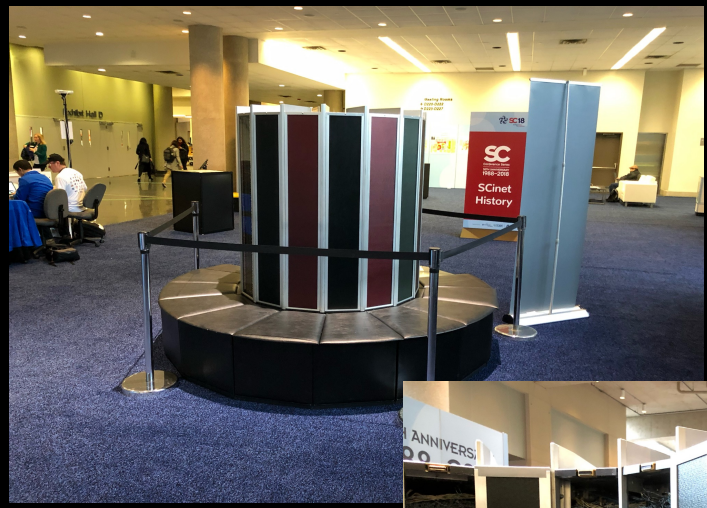

1976

80 MHz
160 MFlops
8 MByte memory
300 MByte disks
120 kW

Exascale on 1 processor: ((2/3)*3*10^8m/s)/(10^18/s)=2*10^-10m = 20 nanometer ~ 100 atoms

Traditional LCF exascale supercomputers under pressure from cloud (HPDF)

Some progress

1997

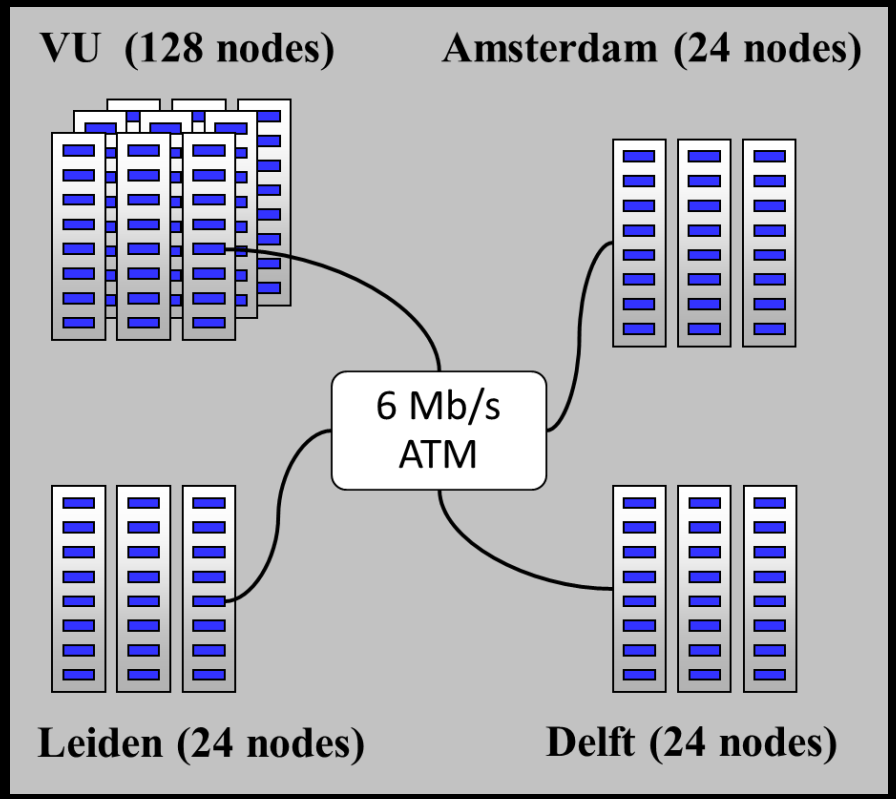

2025

OFC 50
Celebrating 50 Years of Optical Networking and Communications

. . . .

The MacDonalds of Networking, we care about the delicatessen corner

DAS

- No matter what ideas, Henri always wrote a strong proposal
- Solid research at all participants
- Bike Accident, I took over DAS for about a year (thanks Kees Verstoep, all)
- DAS led to:
  - Grid5000 collaboration (Franck Cappello)
  - From Grid5000 website:

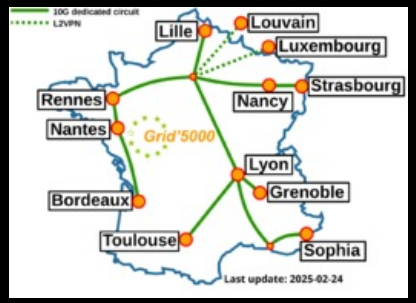

DAS-3

- Netherland testbed composed of 272 nodes (about 800 CPU/cores)
- On the fly network backbone reconfiguration (optical routers with configurable wavelength)
- The software stack is not reconfigurable
- Strong links between DAS-3 and Grid'5000 communities

  - Enabled many international collaborations
  - Inspired Chameleon (Kate Keahey) (private communications)
  - EU SLICES – Research Infrastructure
  - Future Networked Systems

DAS-1 – DAS-6

Advanced School for Computing and Imaging

- Henri had a huge impact on the design of all DAS systems
  - DAS-1: Wide area computing
  - DAS-2: Grid computing
  - DAS-3: Optical Grid
  - DAS-4: Clouds, diversity, green IT
  - DAS-5: Diversity, data-explosion
  - DAS-6: Distributed research ecosystem

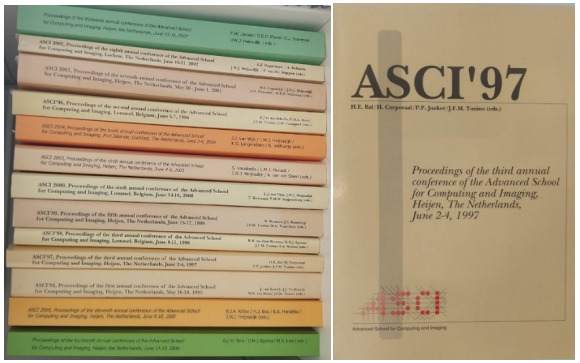

- DAS used by >150 PhD theses

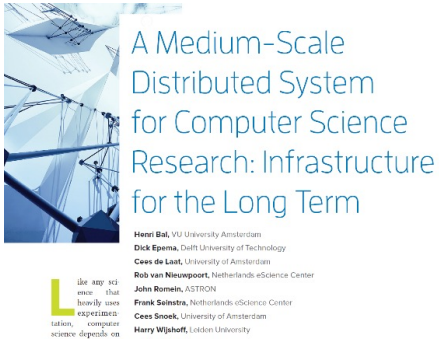

Slide adapted from earlier presentation by Henri ☺


ASCI Conferences
morphed into:
ICT-Open
Compsys

VL-E

- BSIK programme
- Besluit Subsidies Investeringen Kennisinfrastructuur
- 3 relevant programmes
  - Virtual Laboratory for eScience
  - GigaPort
  - Multimedian
- From 2004 till 2008 (extended till end 2009)
- Successfully completed
- Resulted in:
  - National eScience Center
  - BigGrid tier 1
  - COMMIT P20
  - Bridging funding to keep people 4 COMMIT

Project Organization

- Overall project leader L.O.H.
  - Administrative manager J.Vos
  - Daily management Hakam Yakali
- Scientific directorate
  - L.O.H.
  - Heeren/Bouten/Out/van Driel
  - Sloot/van Liere/ Afsarmanesh
  - Bal/de Laat/Groep

Project Organization

- Overall project leader L.O.H.
  - Administrative manager J.Vos
  - Daily management Hakam Yakali
- Scientific directorate
  - L.O.H.
  - Heeren/Bouten/Out/van Driel
  - Sloot/van Liere/ Afsarmanesh
  - Bal/de Laat/Groep

Part. Role	Part. No	Assistant to Contractor No	Participants Short Name	Number of person/years	Personnel Costs (S)	Equipment	Subcontracting	Dissemination and Protection of Knowledge	Consumables	Overhead	Total Cost	Subsidy Amount Requested
	1		WTCW	56,67	3.400.000	1.900.000				1.700.000	7.000.000	5.100.000
	2		NBIC	0,00	0	0				0	0	0
	3		UvA-lvl-CAPS	53,33	3.200.000	800.000				1.600.000	5.600.000	3.200.000
	4		UvA-lvl-AIR	0,00	0	0				0	0	0
	5		UvA-lvl-SCS	23,33	1.400.000	0				700.000	2.100.000	1.000.000
	6		UvA-SWI	11,11	666.667	0				333.333	1.000.000	600.000
	7		UvA-SILS	21,11	1.266.667	500.000				633.333	2.400.000	900.000
	8		UvA-IBED	18,89	1.133.333	0				566.667	1.700.000	800.000
	9		NIKHEF	28,89	1.733.333	100.000				866.667	2.700.000	1.200.000
	10		AMOLF	14,44	866.667	700.000				433.333	2.000.000	800.000
	11		CWI	11,11	666.667	0				333.333	1.000.000	600.000
	12		SARA	5,56	333.333	0				166.667	500.000	300.000
	13		AMC	8,89	533.333	400.000				266.667	1.200.000	400.000
	14		VUA-CS	28,89	1.733.333	200.000				866.667	2.800.000	1.400.000
	15		VUMC	4,44	266.667	400.000				133.333	800.000	200.000
	16		TUD-ITS	17,78	1.066.667	0				533.333	1.600.000	800.000
	17		TNO-TPD	7,22	433.333	0				216.667	650.000	300.000
	18		TNO-Voed.	5,56	333.333	0				166.667	500.000	250.000
	19		WCFS	8,89	533.333	0				266.667	800.000	300.000
	20		Unilever	7,78	466.667	0				233.333	700.000	200.000
	21		ATO	10,00	600.000	0				300.000	900.000	400.000
	22		FCDF	7,78	466.667	0				233.333	700.000	200.000
	23		IBM	8,89	533.333	750.000				266.667	1.550.000	450.000
	24		LogicaCMG	17,78	1.066.667	0				533.333	1.600.000	200.000
	25		Philips	22,22	1.333.333	0				666.667	2.000.000	0
	26		FEI	13,33	800.000	0				400.000	1.200.000	400.000
				413,89	24.833.333	5.750.000	0,00	0,00	0,00	12.416.667	43.000.000	20.000.000

Contribution Johan Vos

VLE middleware

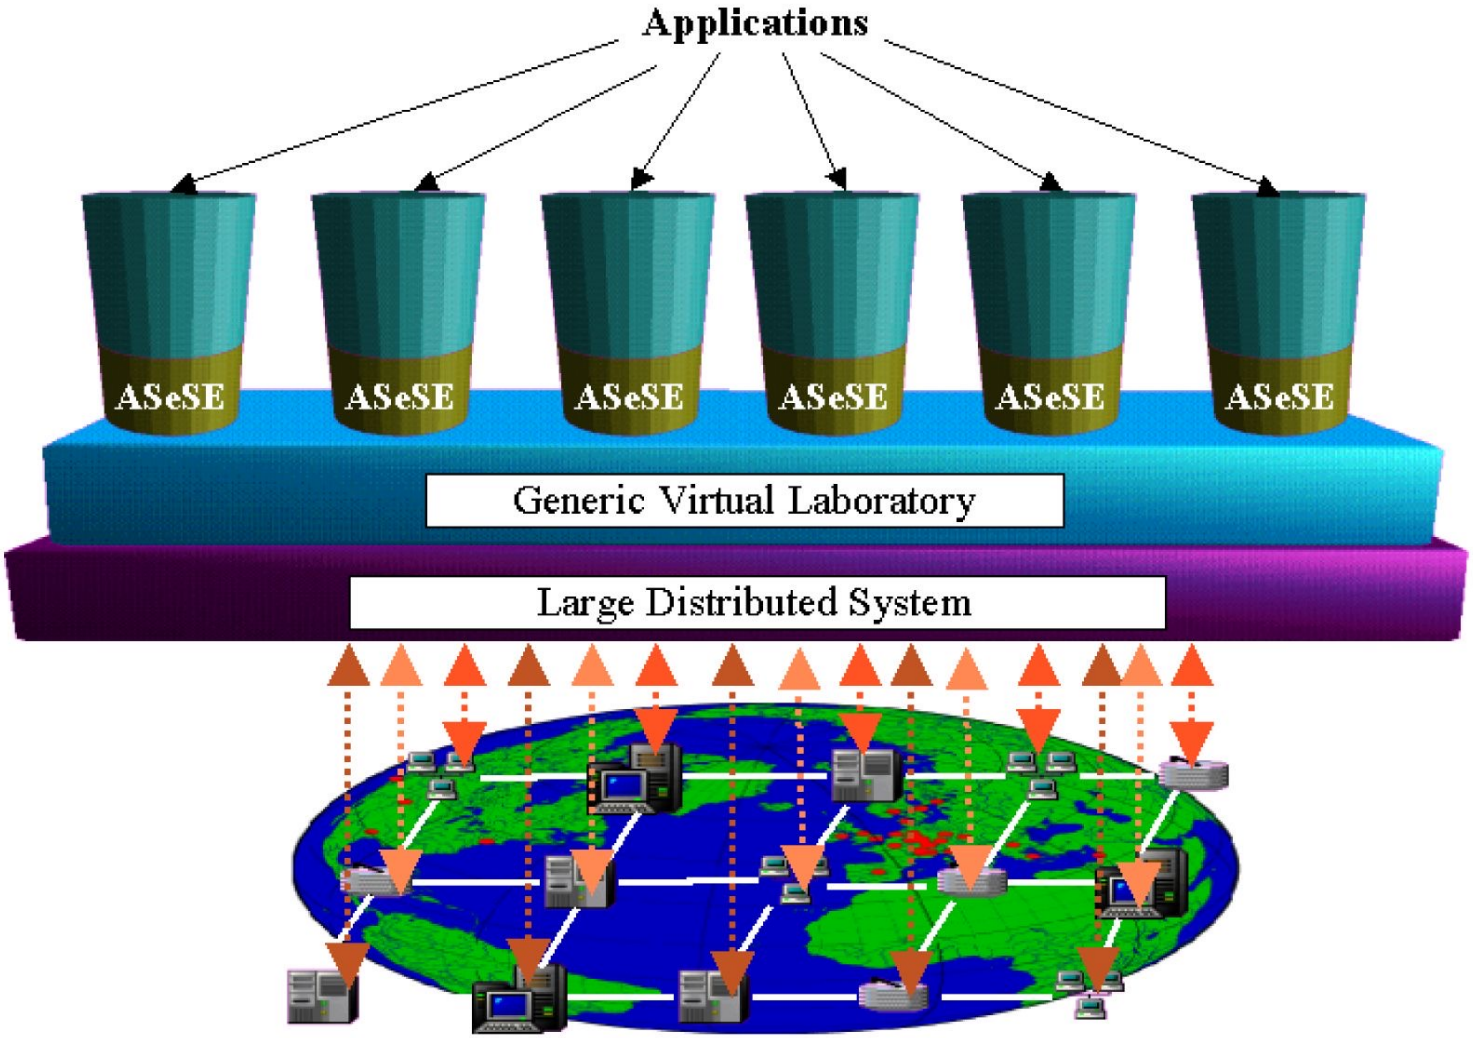

OGSA/OGSI

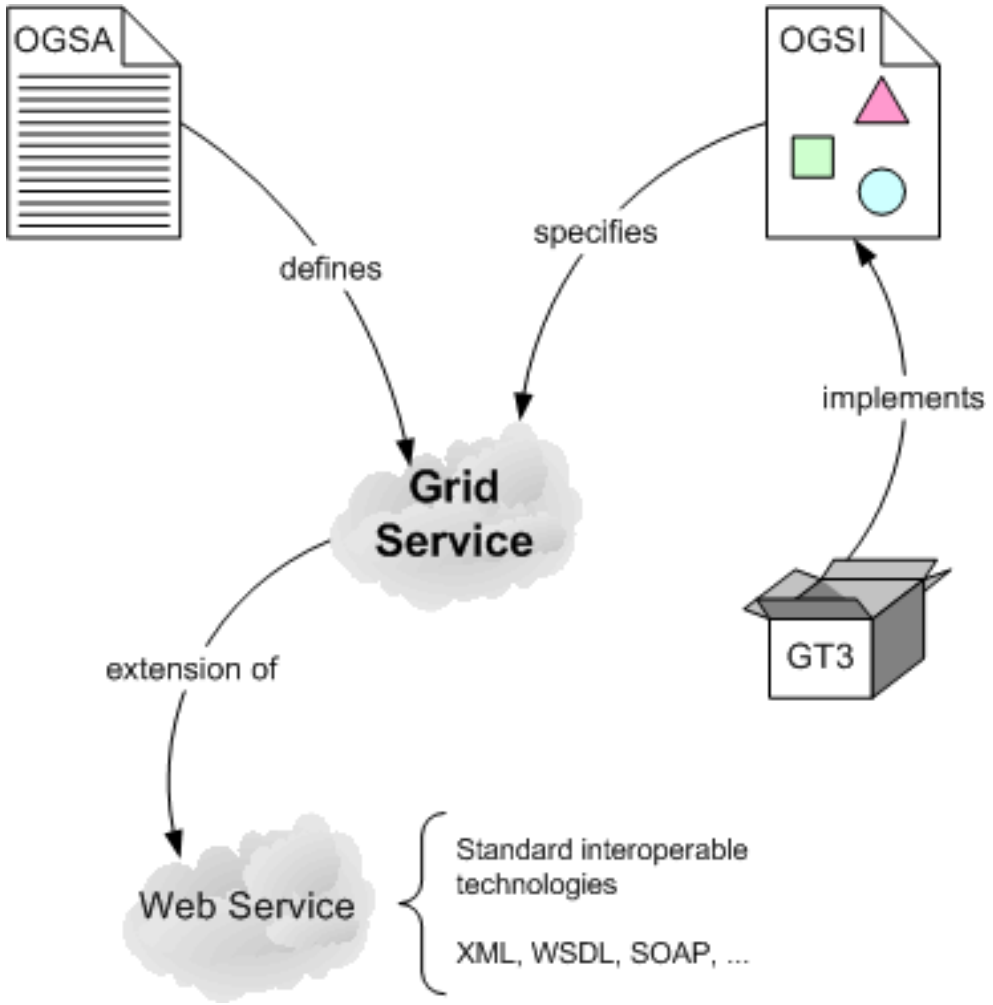

Ref: GT3 tutorial

Research Program

- Large Scale Distributed Systems & Grid
  - Large Scale Distributed Systems (Bal)
  - Grid (Groep/de Laat)
  - Optical networking (de Laat)
  - System for experimentation (Berg, Bal, deLaat,Groep)

VL-e

- 2006
  - Bob Hertzberger retires
  - But not really
  - I became part of VL-e directorate

Contribution Arjen van Rijn

Contribution Arjen van Rijn

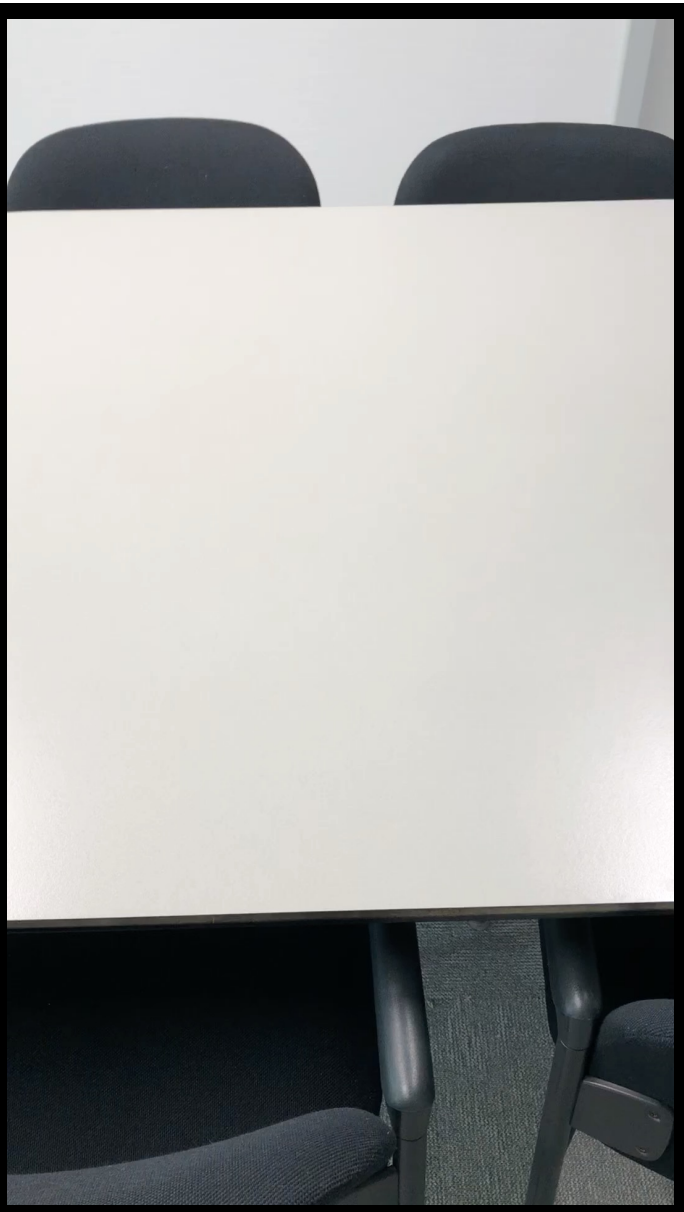

Motto of the VL-e worker

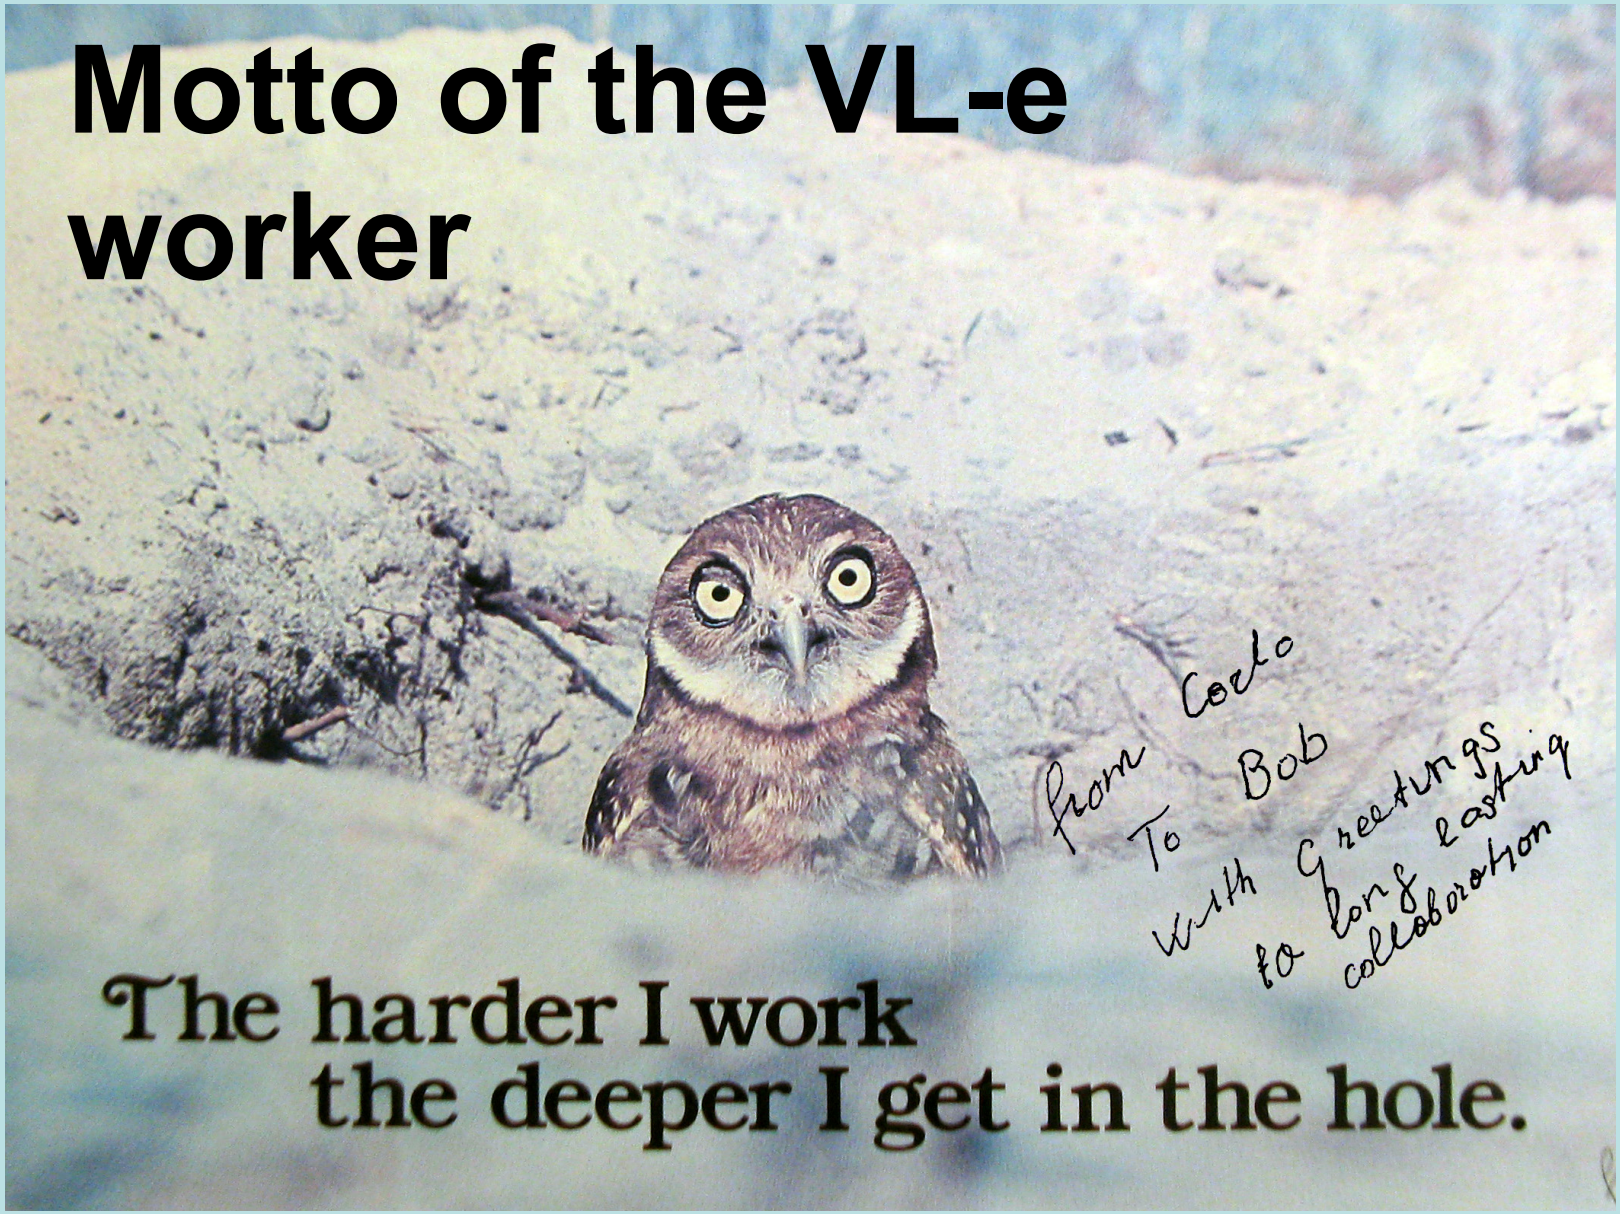

vl-e

COMMIT/

COMMIT P20

P.I. Henri Bal

The project is structured into the following 10 Work Packages:

1. Semantic Description of Infrastructure (SemDIF) – Paola Grosso
2. P20-WP2 Programmable Infrastructure (PIF) – Lydia Meijer
3. Security of (Virtual) e-Science Infrastructure (SeSI) – Cees de Laat
4. Programming Systems for elastic Cloud applications (ProSysClouds) - Henri Bal
5. e-Science applications on large-scale hybrid distributed systems (SALSY) - Frank Seinstra
6. Generic scheduling and co-allocation (SCHOALA) – Dick Epema
7. Application type-specific Scheduling Support (ATYSUP) – Dick Epema
8. Workflow Process Modeling & Management (WOPMOM) – Marian Bubak
9. Workflow Sharing and Reproducability (WSAR) – Adam Belloum
10. Workflow and Application Component Integration (WACI)– Tilo Kielmann

Big Data Sharing use cases placed in airline context

Global Scale

National Scale

City / regional Scale

Campus / Enterprise Scale

Aircraft Component Health Monitoring (Big) Data
NWO **CIMPLO project**
4.5 FTE

Cargo Logistics Data
(C1) DaL4LoD
(C2) Secure scalable policy-enforced distributed data Processing
(using blockchain)

NLIP iShare project

iSHARE
powered by NLIP

Cybersecurity Big Data
NWO COMMIT/
SARNET project
3.5 FTE

COMMIT/

CH

Dutch
Young
VIDI
URC

3 VEN
EuroP
5 Bes

ONVZ
Lotto-
KLM

COMMIT/

TOP PUBLICATIONS

IEEE Computer 2016: strategy of and achievements with the DAS systems
IEEE Computer 2015: datacenter management jointly with industry (Bitbrains)
ICDCS 2015: Scalable, Instant VM Startup for IaaS Clouds
HPDC 2014: storing VM images in clouds
SIGMETRICS 2014: big data processing with MapReduce
ICDCS 2014: distributed middleware for MapReduce and stream processing
TPDS 2014: Parallel Workload Modeling with Realistic Characteristics
ISWC 2013: distributed reasoning on dynamic semantic web data
SC 2013: VM deployment
SC 2013: portfolio scheduling
IEEE Big Data 2013: first Dutch article in a top big data conference
IEEE Internet Computing 2013: Distributed Computing on an Ensemble of Browsers
IEEE Internet Computing 2012: Enabling Web Services to Consume and Produce Large Datasets
TPDS 2012: Cost-driven Scheduling of Grid Workflows Using Partial Critical Paths
IEEE Internet Computing 2011: Grid Computing Workloads
ISWC 2011: an OWL reasoner that scales to 1 billion triples
TPDS 2011: performance analysis of cloud services, this work is the **highest cited** in the highest impact journal in the field, for the period 2009-2014 (source: Google Scholar)
HPDC 2011: Incremental Placement of Interactive Perception Applications

Contribution
Arnold Smeulders

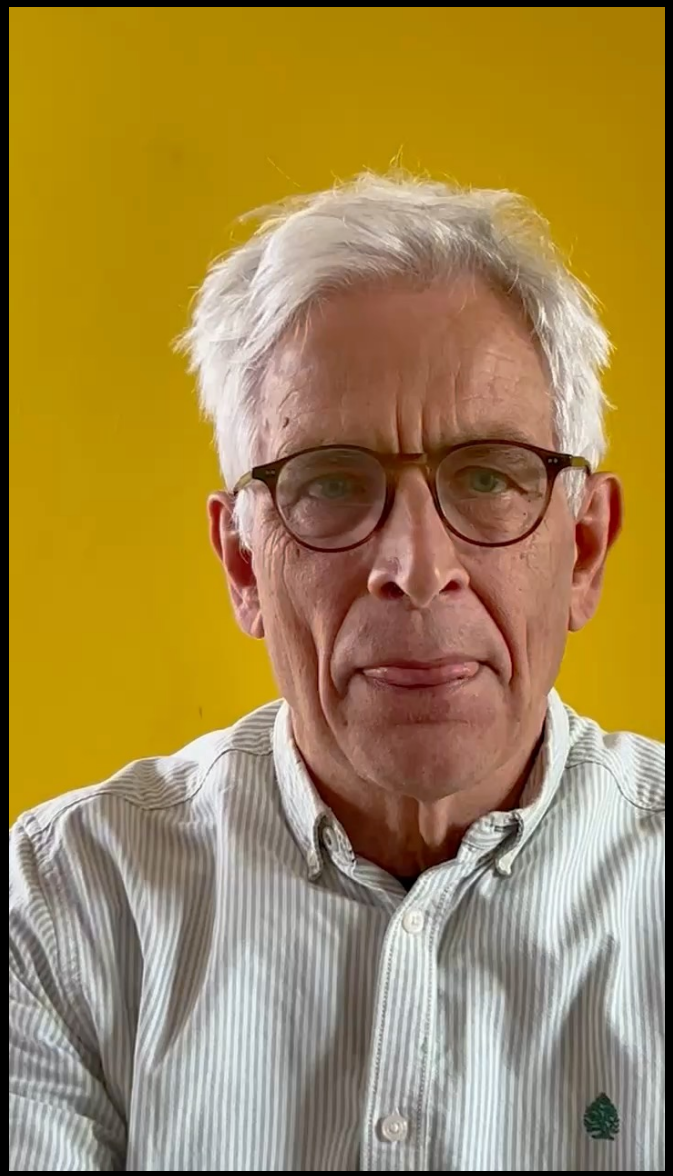

Arnold Smeulders: The Grandmaster of AI

YES

NO

Contribution Bob Hertzberger

Contribution Bob Hertzberger

Complementary skills

	Henri		
Diplomatic			
Approachable			
Management	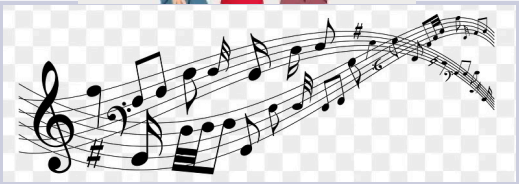		
Methodology	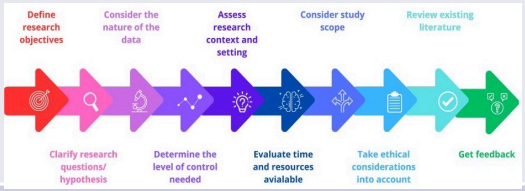		
Field	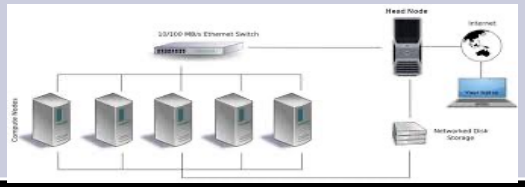		

Complementary skills

	Henri	Cees	
Diplomatic			
Approachable			
Management	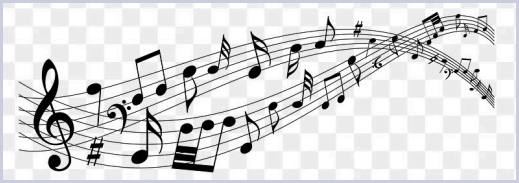		
Methodology	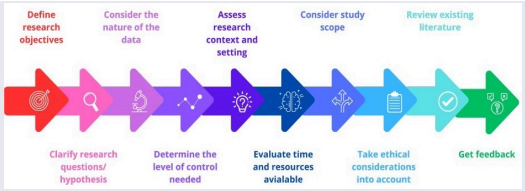		
Field	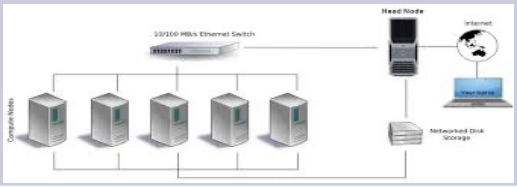	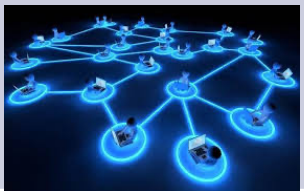	

Complementary skills

	Henri	Cees	Bob
Diplomatic			
Approachable			
Management	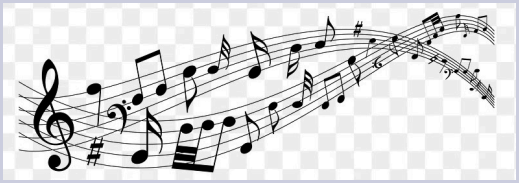		
Methodology	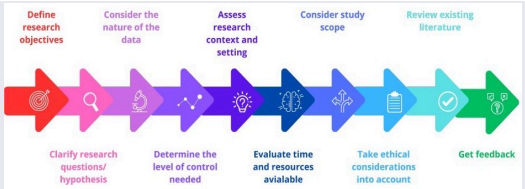		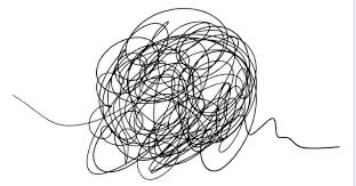
Field	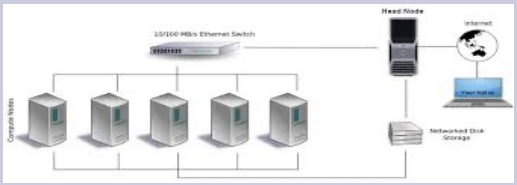	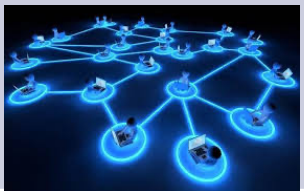	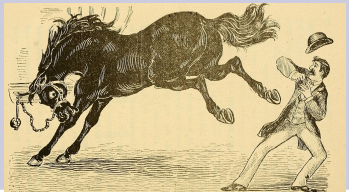

From

Infinite Do Loops

to

Never Work Again

On the occasion of the retirement of

Prof.dr.ir. Henri Bal

by

Prof.dr.ir. Cees de Laat

Finally,

- Henri has had a tremendous impact on computer science in the Netherlands
- Henri has performed world class fundamental science
- He was the silent force enabling many different large projects
- A very pleasant person to work with! Never got him angry
- Many thanks for ~30 years of inspiringly working together!

Finally.