Computing for the Future of the Planet

Andy Hopper

The Computer Laboratory
University of Cambridge

Optimal Digital Infrastructure

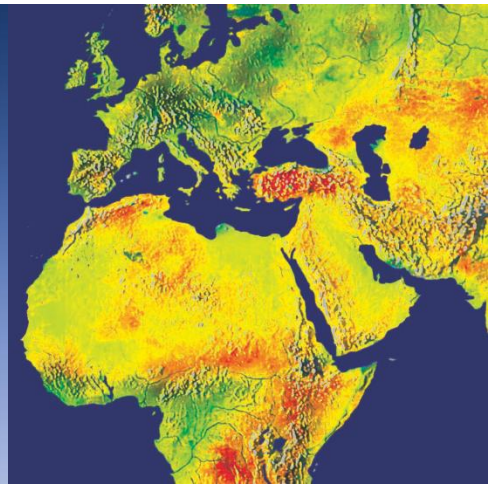

Sense and Optimise

Predict React Automate

Physical to Digital

Andy Hopper and Andrew Rice, Computing for the Future of the Planet, *Phil. Trans. R. Soc. A* 366(1881):3685–3697, 2008
&
https://www.cl.cam.ac.uk/research/dtg/www/research/

Framework

1. Optimal Digital Infrastructure – “Green Computing”
2. Sense and Optimise – “Computing for Green”
3. Predict, React, Automate - “Guaranteed Computing”
4. Digital Alternatives to Physical Activities – “Physical to Digital”

Green Computing

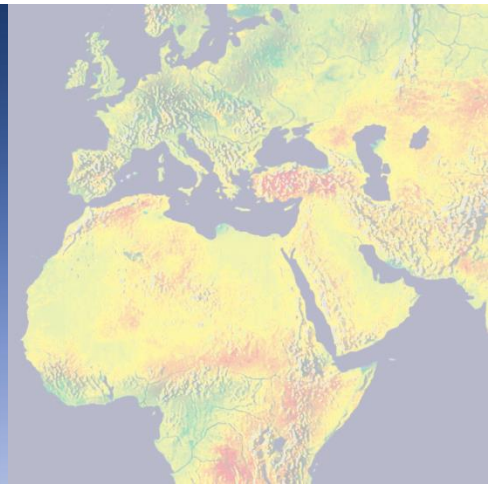

Computing for Green

Guaranteed Computing

Physical to Digital

1 – Green Computing

- Data centre design
- Energy-proportional processing, networking, storage
- Direct use of renewable energy
- Use of surplus renewable energy
- Use of very remote renewable energy?

Sun

Energy-proportional processing

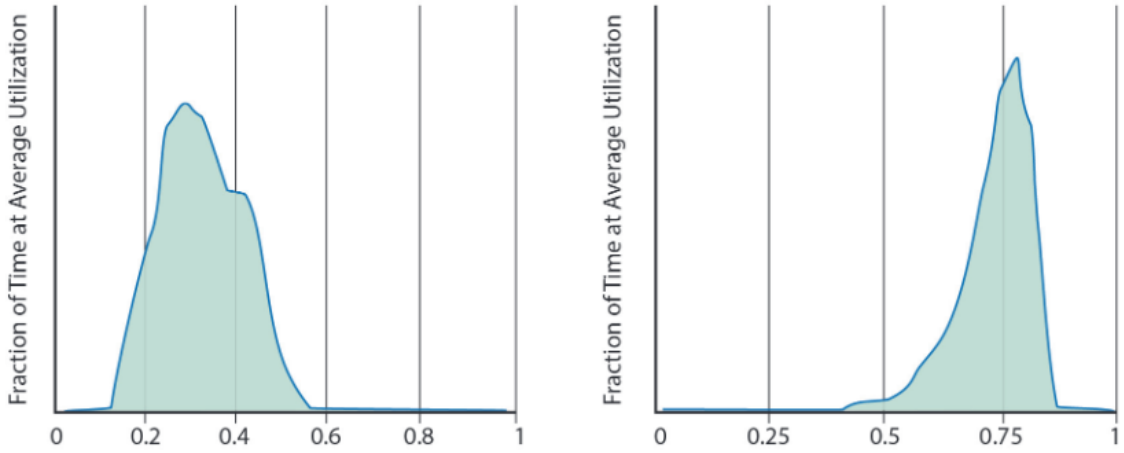

Figure 5.4: Average activity distribution of a sample of two Google clusters, each containing over 20,000 servers, over a period of 3 months (January-March 2013).

The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines, Luiz André Barroso, Jimmy Clidaras, Luiz André Barroso

Energy-proportional storage (Facebook)

- Hot storage: read a lot
  - Replication/striping across data centres
  - 3.5x stretch factor for resilience and performance
  - Data cools by 500x over 1st year
- Warm storage: read a bit
  - 2.5x stretch factor at lower performance
  - Migrate 1 month old data
- Cold storage: read almost never
  - Cheap drives by the million powered 1/15 of the time
  - 1.4x stretch
- Very cold storage: read never?
  - Optical (eg Blu-ray) thus little energy storage cost

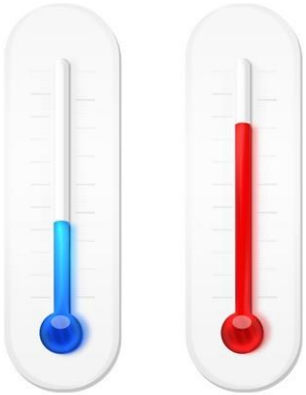

Use of Renewable Energy

- Execute computing tasks where renewable energy is available
- Use renewable energy that cannot be used for another purpose
- Do we ship or replicate programs, data sets, or both?
- Will this change as photonic networks evolve?

Siemens press picture

Technology Perspective

- Technology trends
  - Better hardware performance and power scalability
  - “Internet of Things”
- Workload trends
  - “Big Data”
  - Increasing large data and batch computations
- The future
  - Energy proportional computing, networking, and storage?
  - How much of the world’s computing will use surplus energy?
  - Ultimate processing architecture: latency or energy driven?
  - Ultimate storage architecture: store everything for ever?

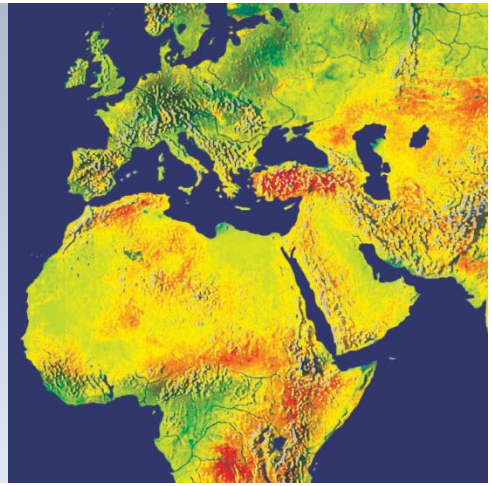

Green Computing

Computing for Green

Guaranteed Computing

Physical to Digital

2 – Computing for Green

- A sensor-based digital model of everything
- “Googling” Earth in space and time
- How do we do it?
  - coverage
  - fidelity
  - scalability
  - performance
  - usefulness

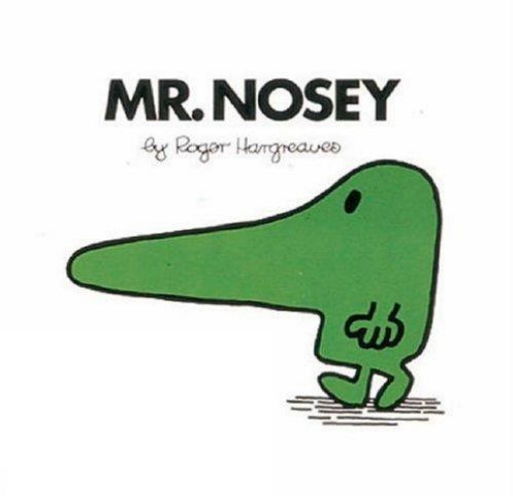

Future Street View – Heat Sensing?

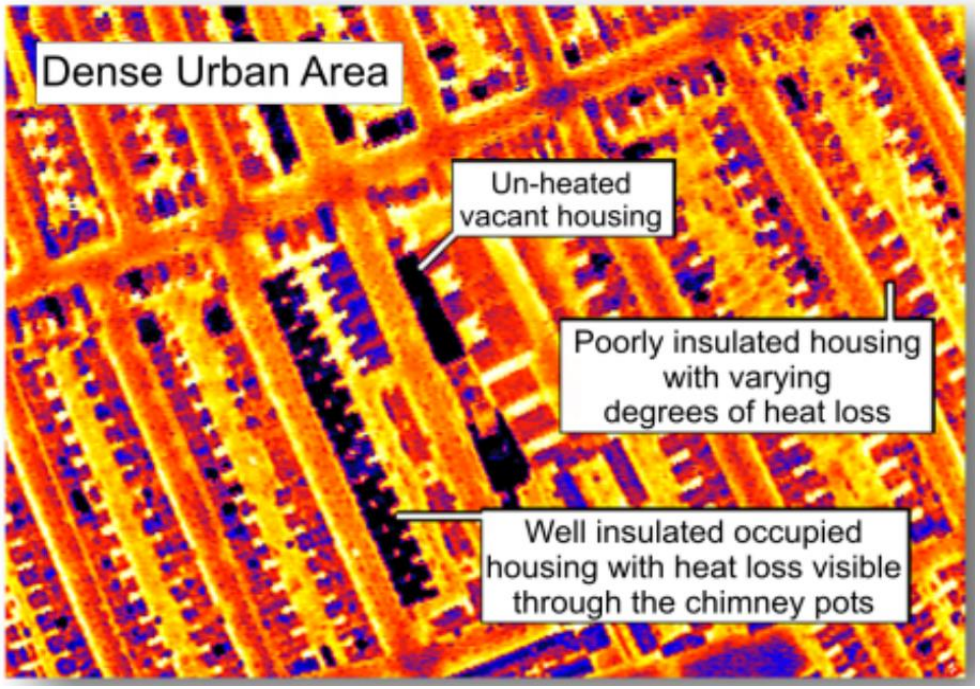

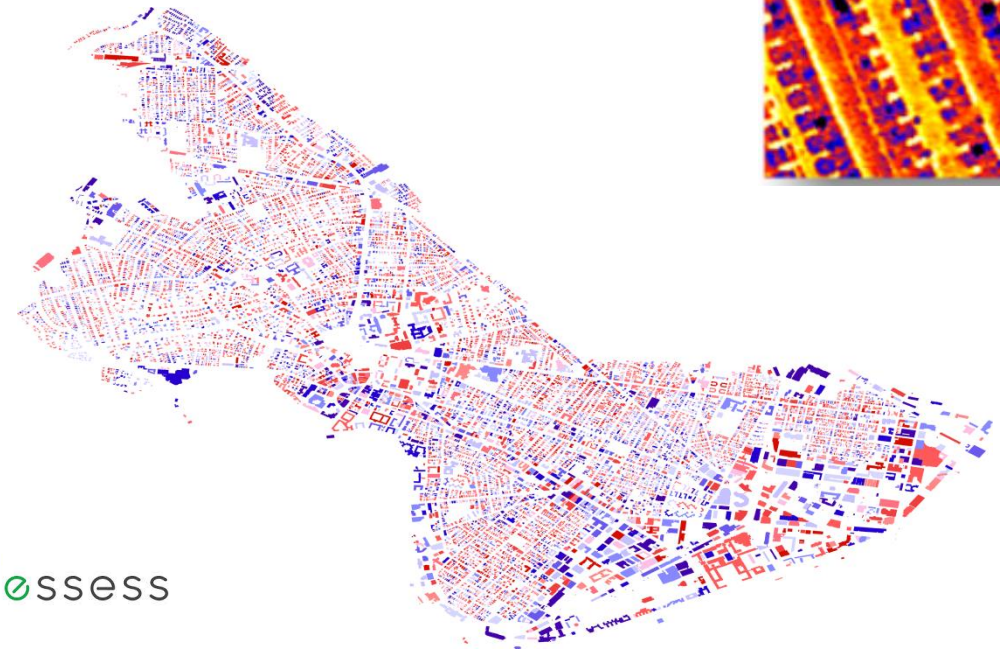

essess

Indoor Location

Sensor

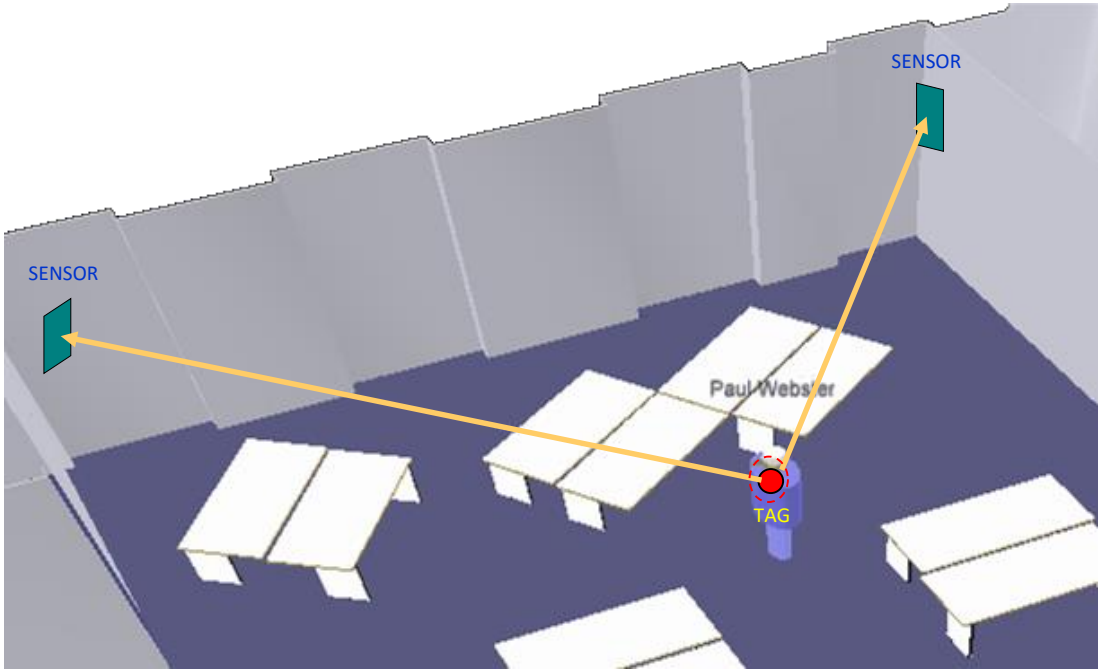

Tag

- Active tag and receiver infrastructure
- 3D accuracy up to 15cm (95% confidence level)
- Max range up to 160m

Ubisense

Tag integrated with tool

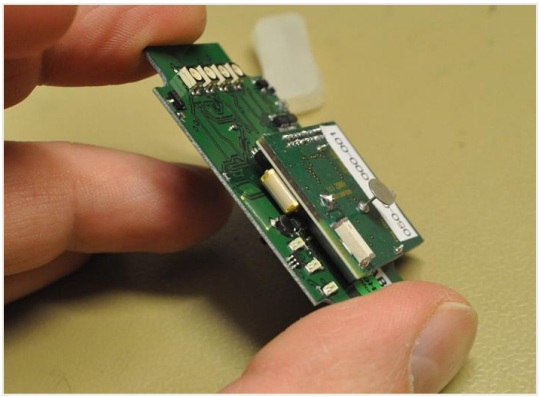

BMW Car Plant, Germany

(Final Assembly Tool Assistance)

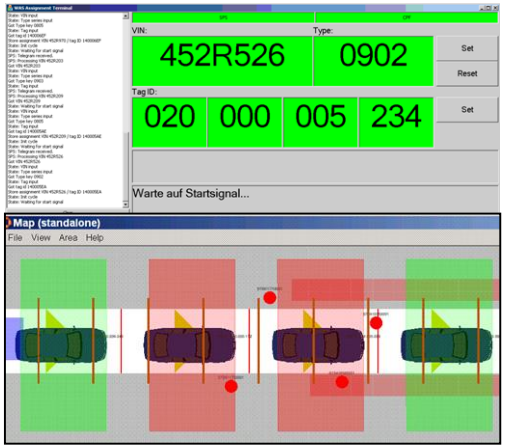

- Tracking tools on a complex production line
- Automatically programs tool with correct settings for each car
- Fewer manual processes, reducing errors

Multiple Airbus sites in Europe

(Process Tracking)

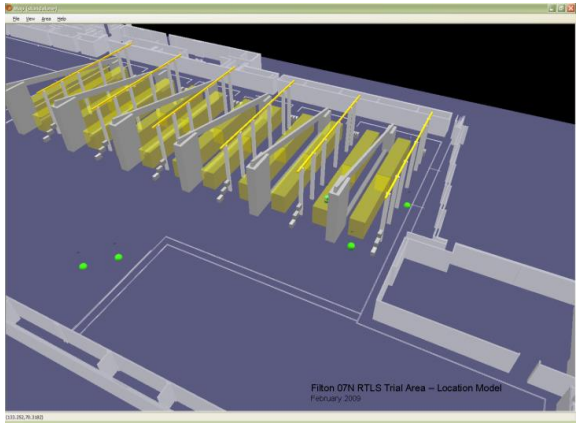

- Process monitoring across multiple factory sites in UK, France, Germany
- Provides central overview of process state
- Automatically updates planning system

Smartphone tracking

R. Harle et al

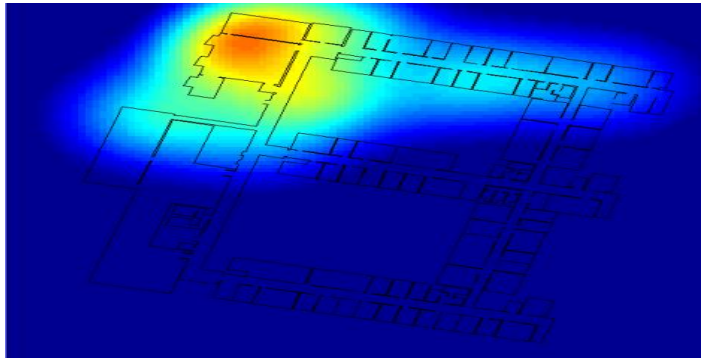

Accelerometers
Gyroscopes
Compass
Barometer
WiFi
Cellular
Bluetooth
NFC, RFID
Whitespace radio
Light
Temperature
etc

Knowledge of environment
Motion models

Fuse

Position

DeviceAnalyzer for Android

Daniel T. Wagner, Andrew Rice and Alastair R. Beresford. Device Analyzer. In HOTMOBILE 2011 12th Workshop on Mobile Computing Systems and Applications, Mar 2011

Global Personal Energy Meter - PEM

A. Rice et al

- Complete
  - all energy accounted for: sensed, embedded, shared, hypothecated
- Accurate / Bounded / Personalised
  - my actions relate to me only
- Sensible
  - incentives work correctly
- Assured
  - rules are understood
  - fidelity / error bounds
  - security / privacy

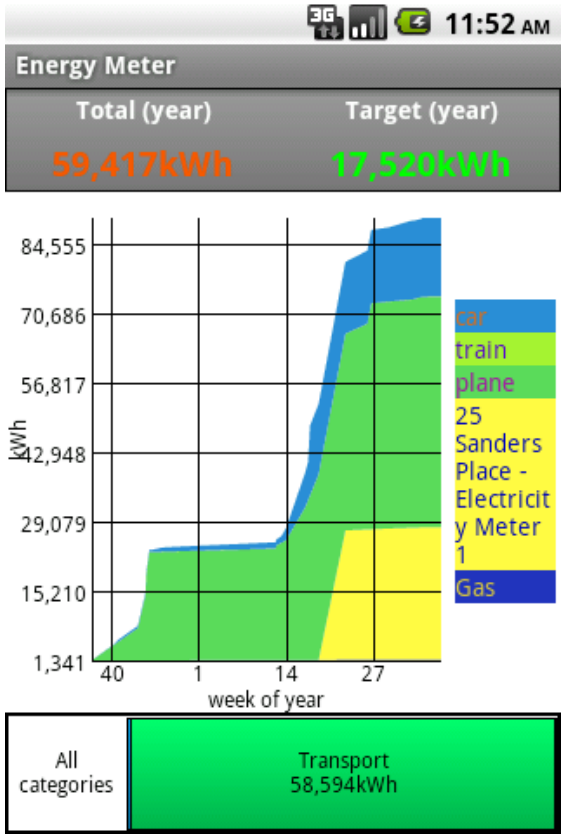

Privacy Dilemmas

- Privacy vs Sustainability
- Privacy vs Public Good
- Privacy vs Unexpected Consequences (eg Facebook)
- Privacy vs Wealth Creation
- Who is Big Brother anyway?

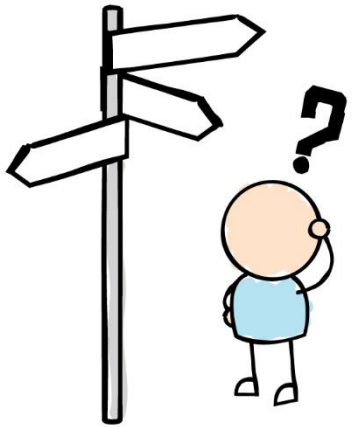

Framework for Big Data

- Collection
  - Transparency, Consent, Purpose, Access, Withdrawal, Accountability

- Governance
  - Who owns it
  - Who do we trust
  - How does business work
  - How does society work

Green Computing

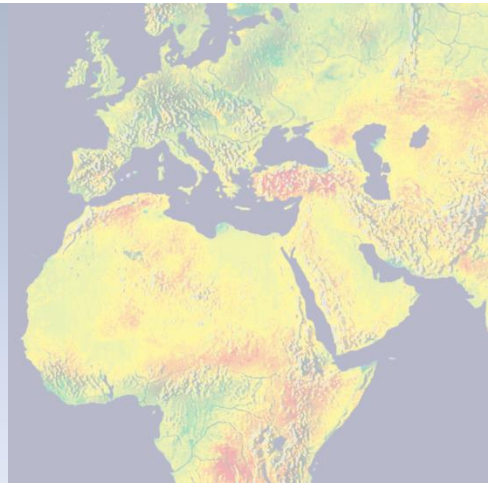

Computing for Green

Guaranteed Computing

Wealth Creation

3 – Guaranteed Computing

- Increasing societal dependence on computing and automation
  - Complex systems
  - Advances in machine learning and artificial intelligence

- Technical challenges
  - correctness, bounding of errors
  - data archiving, code archiving, audit trails, transparency, repeatability, provenance
  - security, privacy

Provenance in Guaranteed Computing Systems

R. Sohan et al

- Complete history of a piece of data and its transformations
  - Can be applied at various levels
- A component of Guaranteed Computing Systems
  - Promulgate all changes to data
  - Invalidate backwards and compute forwards
  - Automatic or on-demand implementation
- Uses
  - Validation and Reproducibility
  - Audit and Compliance
  - System Optimisations
  - Reversible consent for exploratory use of Big Data?

Ongoing Projects

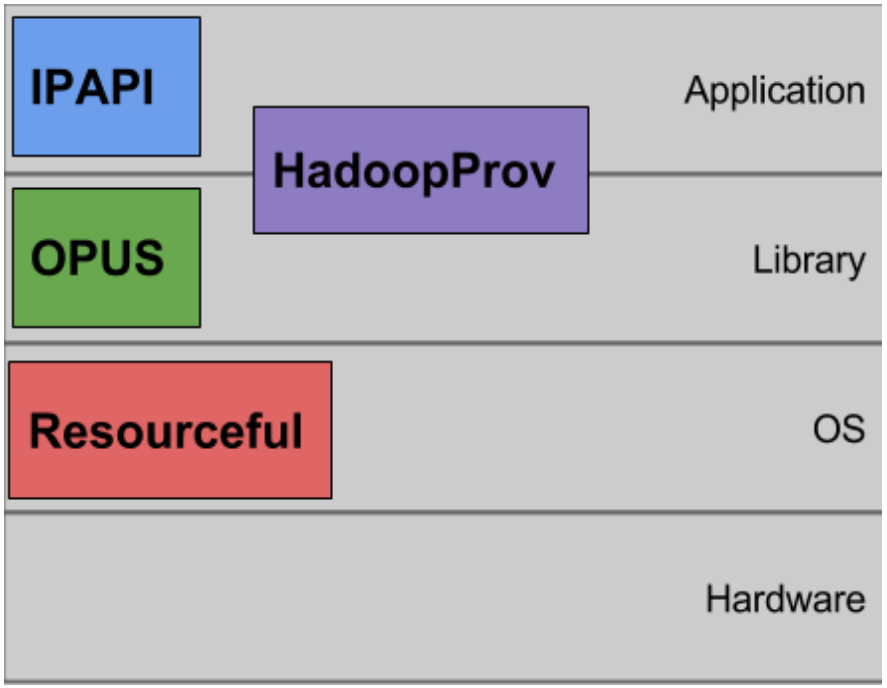

IPAPI – An Improved Provenance API

HadoopProv – A tool for augmenting “Big Data” programming

OPUS – Observational Provenance in User Space

Resourceful – System-call level Resource Accounting

HadoopProv: Provenance For MapReduce

S. Akoush et al

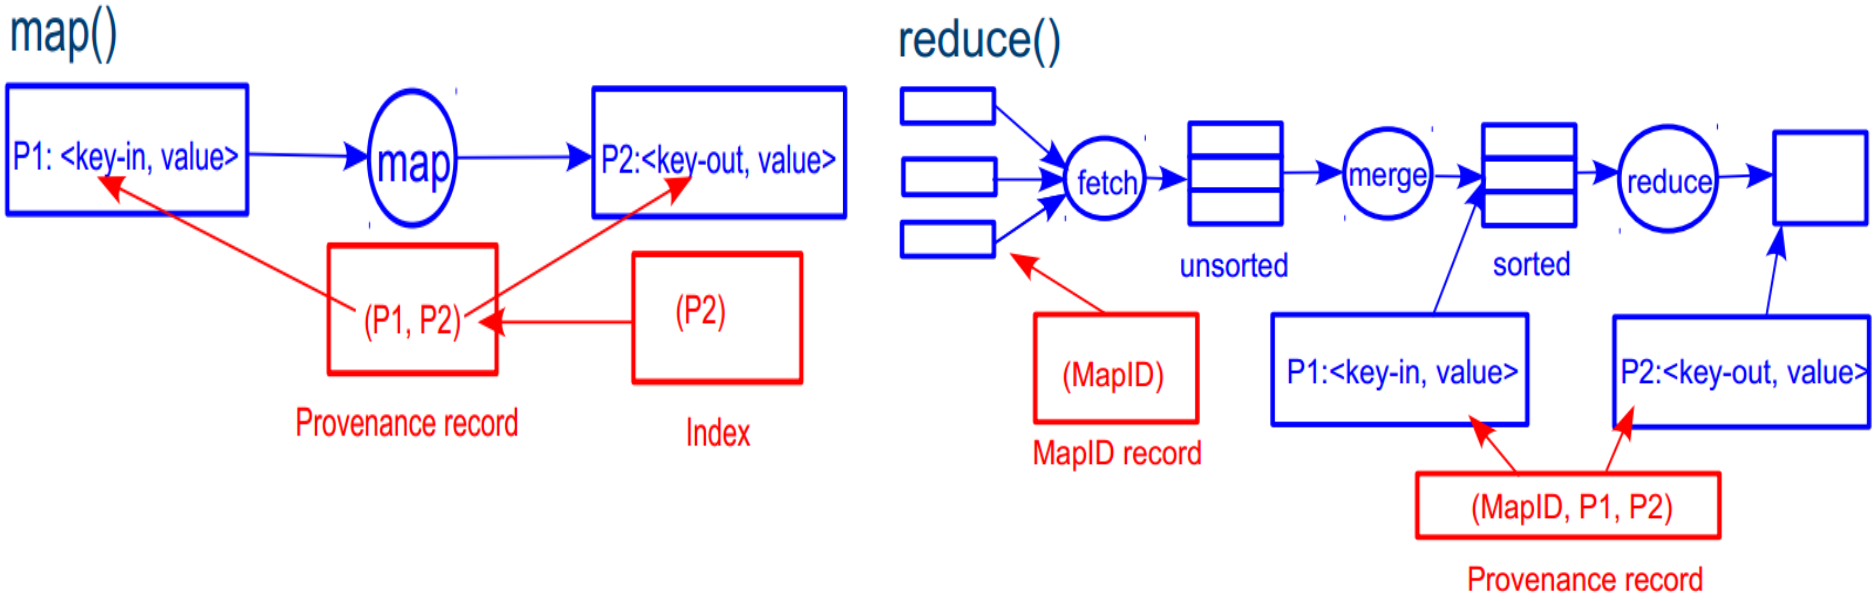

- Capture key-value record dependency for across entire work-flow
- Low (<10%) temporal overhead
- Enables forward and backward key-value trace
- Uses: verification, validation, subset processing

OPUS: Observational Provenance In User-Space

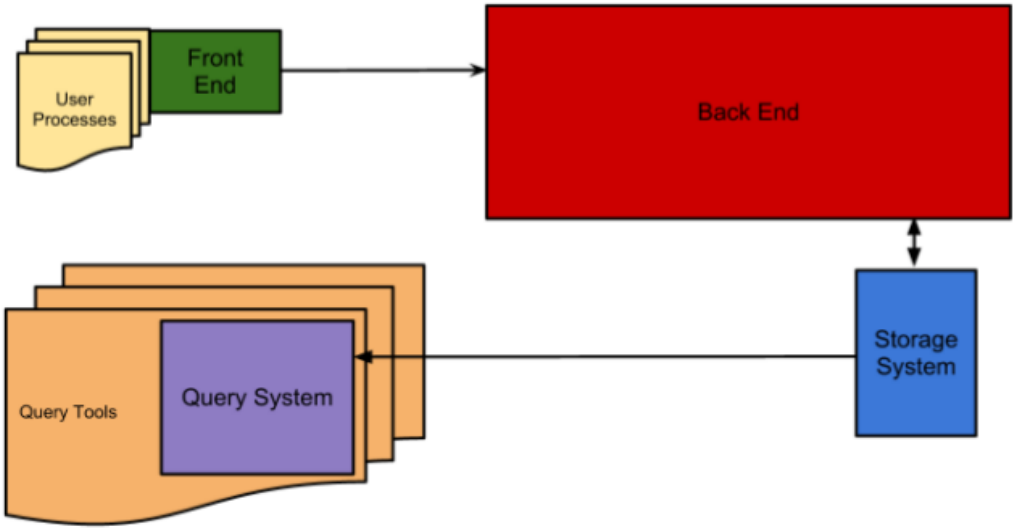

- General purpose process-level provenance capture for POSIX
- Record all process I/O calls at library level
- Low use barrier, low overhead, always on

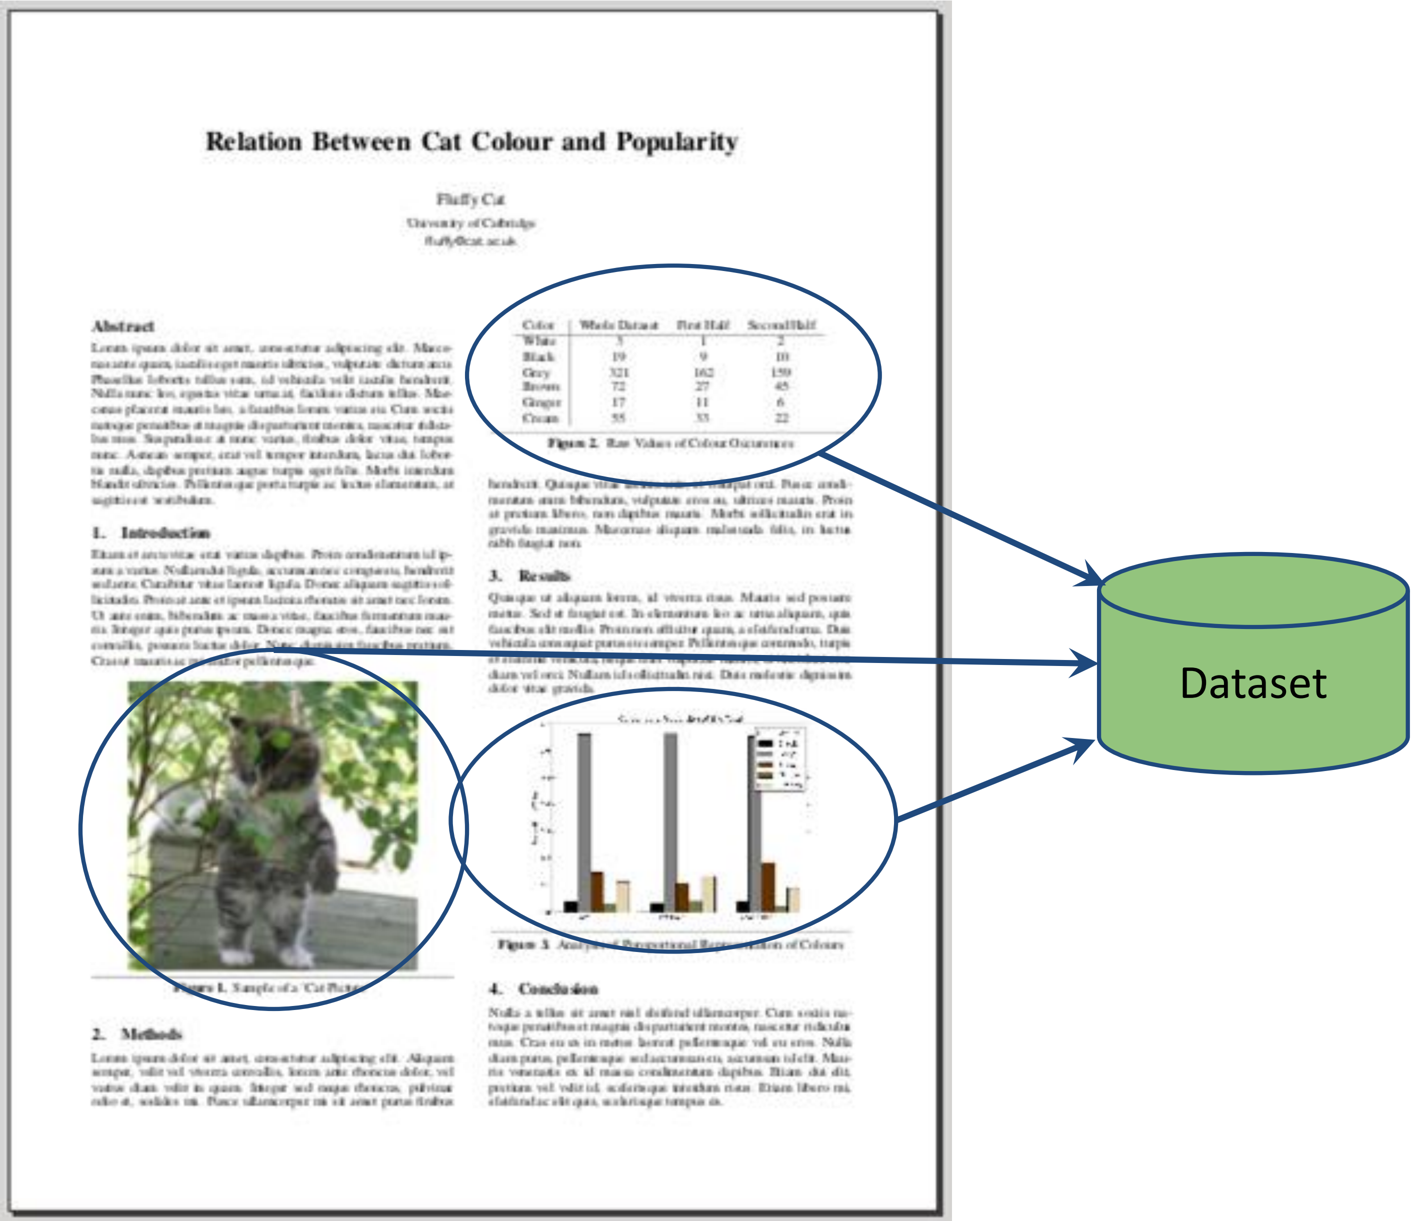
Relation Between Cat Colour and Popularity
Abstract
1. Introduction
2. Methods
3. Results
4. Conclusion
Dataset

How?

When?

Why?

Green Computing

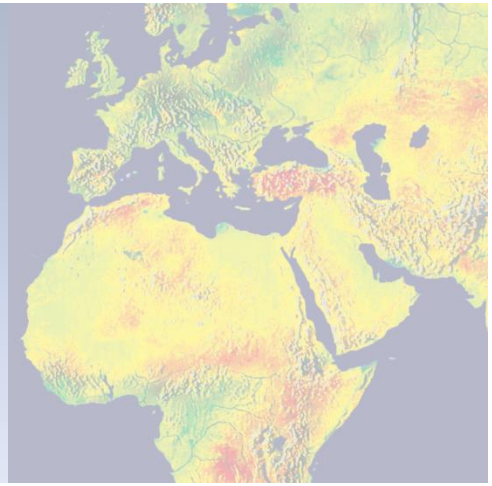

Computing for Green

Guaranteed Computing

Wealth Creation

4 – Wealth in Cyberspace?

- Can we construct a digital world in which we can conduct our lives?
  - On a ultra-cheap platform
  - Using miniscule power
  - Fed with sensor data from the real-world
  - Accessible to everyone and everything

- Opportunity for the “Developing World” to bypass intermediate steps
  - Create wealth with a smaller footprint

Mobile traffic in Rwanda

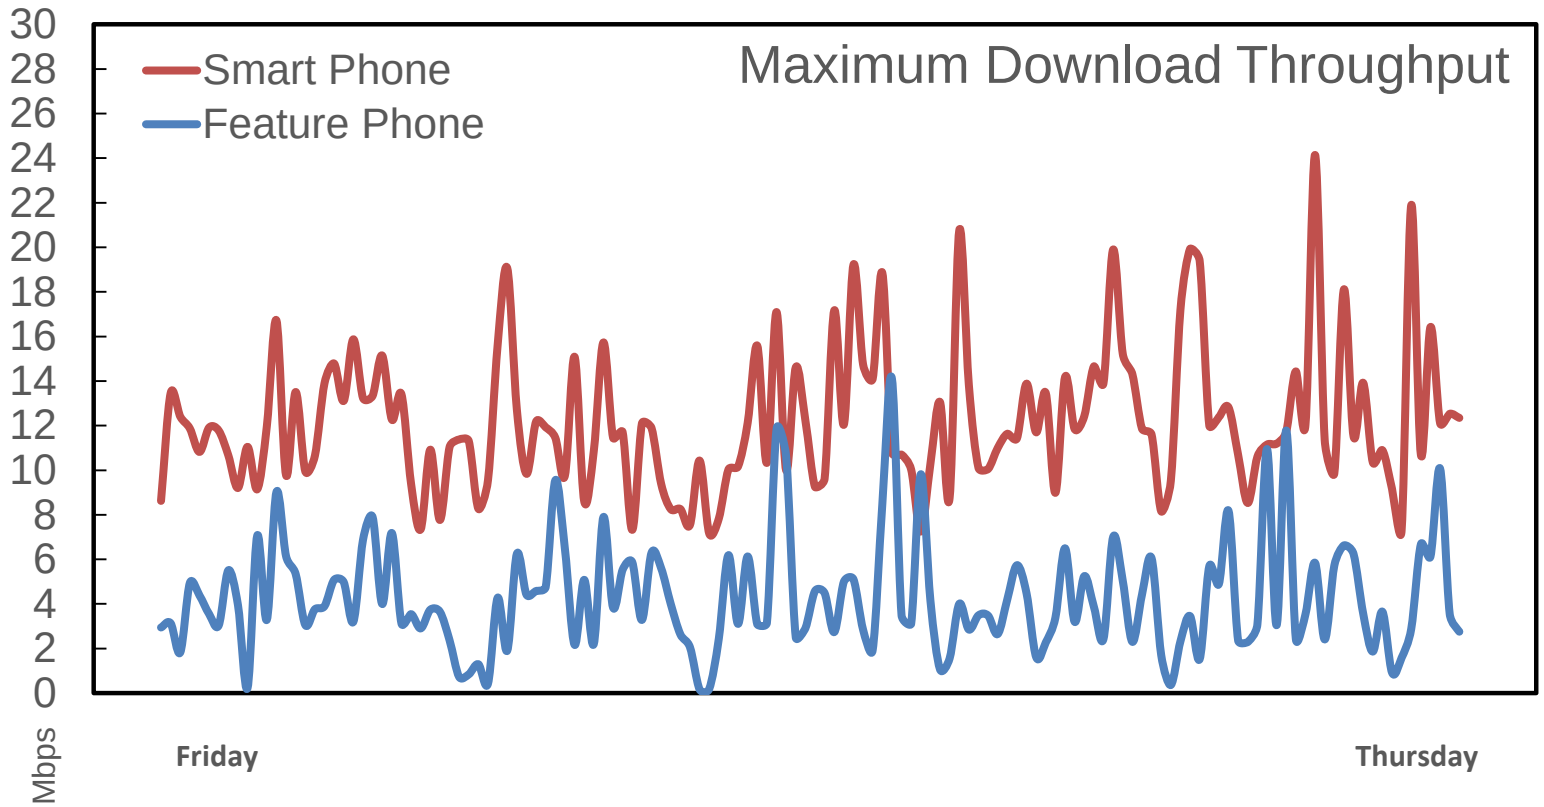

Its happening!

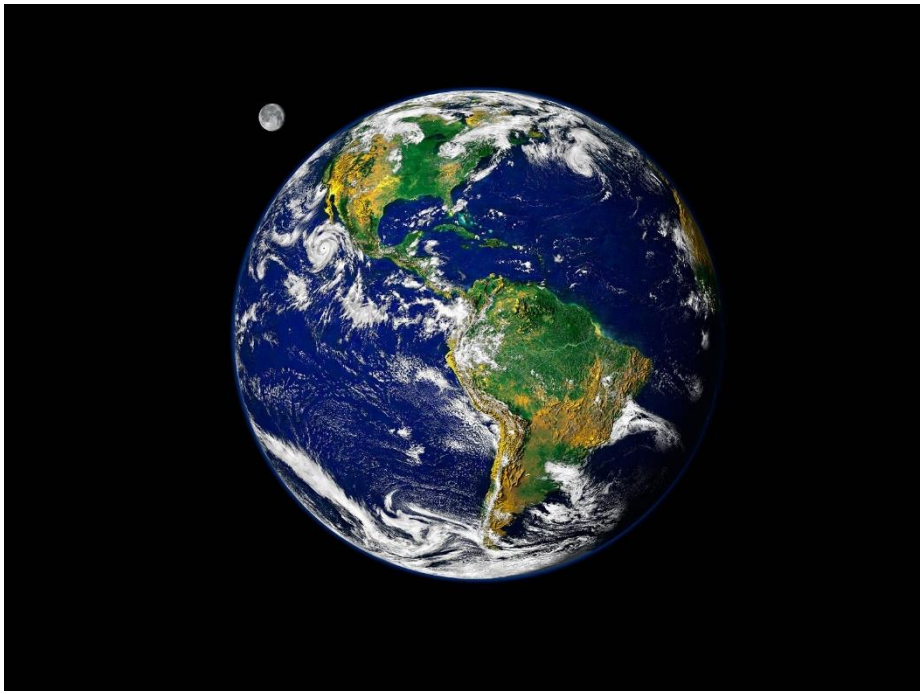

- Coverage of cellular infrastructure is over 75%
- Auto-rickshaw: mGaadi, PoochO
- Amazon Mechanical Turk is a crowdsourcing marketplace for work
- MOOCs for education, eg: isaacphysics.org
- etc

Optimal Digital Infrastructure

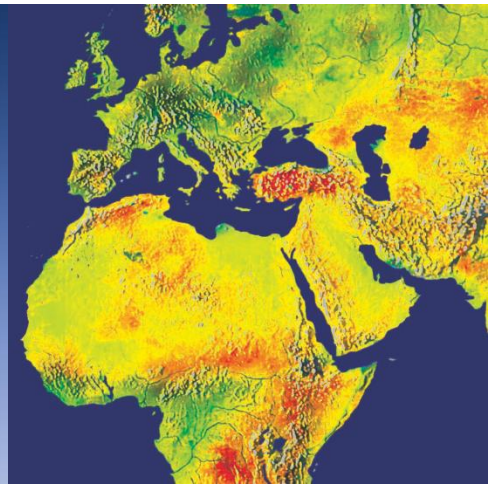

Sense and Optimise

Predict React Automate

Physical to Digital

Andy Hopper and Andrew Rice, Computing for the Future of the Planet, *Phil. Trans. R. Soc. A* 366(1881):3685–3697, 2008
&
https://www.cl.cam.ac.uk/research/dtg/www/research/