From Atoms to Clouds
Ongoing CERCS-Intel Collaborations

Ada Gavrilovska, Karsten Schwan,
PhD & MS students
Matthew Wolf & Chad Huneycutt

CERCS – Intel Interactions

- Multiple, across many faculty and many research and product groups at Intel
  - e.g., HVM project
  - many others – LRB, medical imaging algorithms for next gen hardware, architecture research, multicore curriculum…
- This presentation
  - Atom-based projects – with Intel's Embedded and Communications Group (ECG)
  - Cloud-based efforts – through OpenCirrus
  - (significant cross-leverage across all of these…)

On Clouds: OpenCirrus Update

- OpenCirrus – open cloud platform
  - enabling cloud research and innovation
  - Intel, Yahoo!, HP partnership initiative
  - CMU, Illinois; Korea, Singapore, Malaysia; Germany, Russia;
  - 1000 core membership "fee"
- CERCS/GT is in final stages of formally joining
  - Intel donation of hardware resources
    - 400 + 2 future generation updates
  - Legal requirements non-trivial

CERCS/GT Contributions

- Provide regional capabilities
- vManage/Monalytics management infrastructure (joint with HP)
- Diverse workloads
  - HPC, large-scale SOA applications, etc…
- Focus on heterogeneity
  - advanced next generation platforms, presence of accelerators…
- Leverage on-going GreenIT effort
  - measurement-based approach to energy-aware management
- Leverage Georgia Tech connectivity for exploration of geographically distributed cloud opportunities
  - initiated with CERCS-Emory collaboration

Current Status

- Awaiting arrival of first hardware increment (& final legal signatures)
- Evaluation of basic cloud infrastructure
  - Eucalyptus (used at HP)
  - Tashi and Zoni (Intel)
  - XCP – in initial release
- Expected to start providing services late summer/fall
  - initial plans for split mode:
    - existing open source cloud stack
    - advanced vManage management stack

On Atoms: CERCS-ECG

- Relationship seeded in IXP work
- IXP uE cores part of Intel's embedded Tolapai SoC – a tightly coupled HVM platform
- Shift of vStore – our joint effort with Motorola – to Atom-based platforms
- vStore evolves in a larger cloud@home project
- Atom-Nehalem asymmetric multicore prototype

cloud@home

- @home, in-Vehicle, office… personal infrastructures increasingly more computationally rich
  - Rich resources == rich services
    - e.g., personalized entertainment, …
  - Interactive, time-sensitive, sensing-based tasks
    - e.g., security camera analysis, …
- Opportunities to tap into aggregate infrastructure resources and enrich end-user service experiences
  - Benefit from various service interface modalities
    - e.g., migrate between handheld and large flat screen…
- Key challenge is to deal with, and benefit from, resource heterogeneity
- Interface to open/public clouds
  - e.g., EC2, Facebook, …

cloud@home Approach

- Create ‘compute clouds’ in @home environments
- Leverage platform-level virtualization technology
  - dynamic deployment, migration, isolation…
- Combine with ‘service-level’ virtualization
  - services provided by infrastructure hardware components
    - storage, network, display, inputs, …
    - graphics, crypto, compression, …
  - ability to adapt service to best utilize underlying physical resources
  - ability to dynamically ‘compose’ platforms with required capabilities
- Enable infrastructure-wide management policies
  - e.g., battery life, application performance, platform efficiency…

Why Cloud @home

- Rich resources == rich services
  - e.g., personalized entertainment, …
- Benefit from various service interface modalities
  - e.g., migrate between handheld and large flat screen…
- Interactive, time-sensitive, sensing-based tasks
  - e.g., security camera analysis, …
- Privacy concerns

- Interface to open/public clouds
  - e.g., EC2, Facebook, …

cloud@home

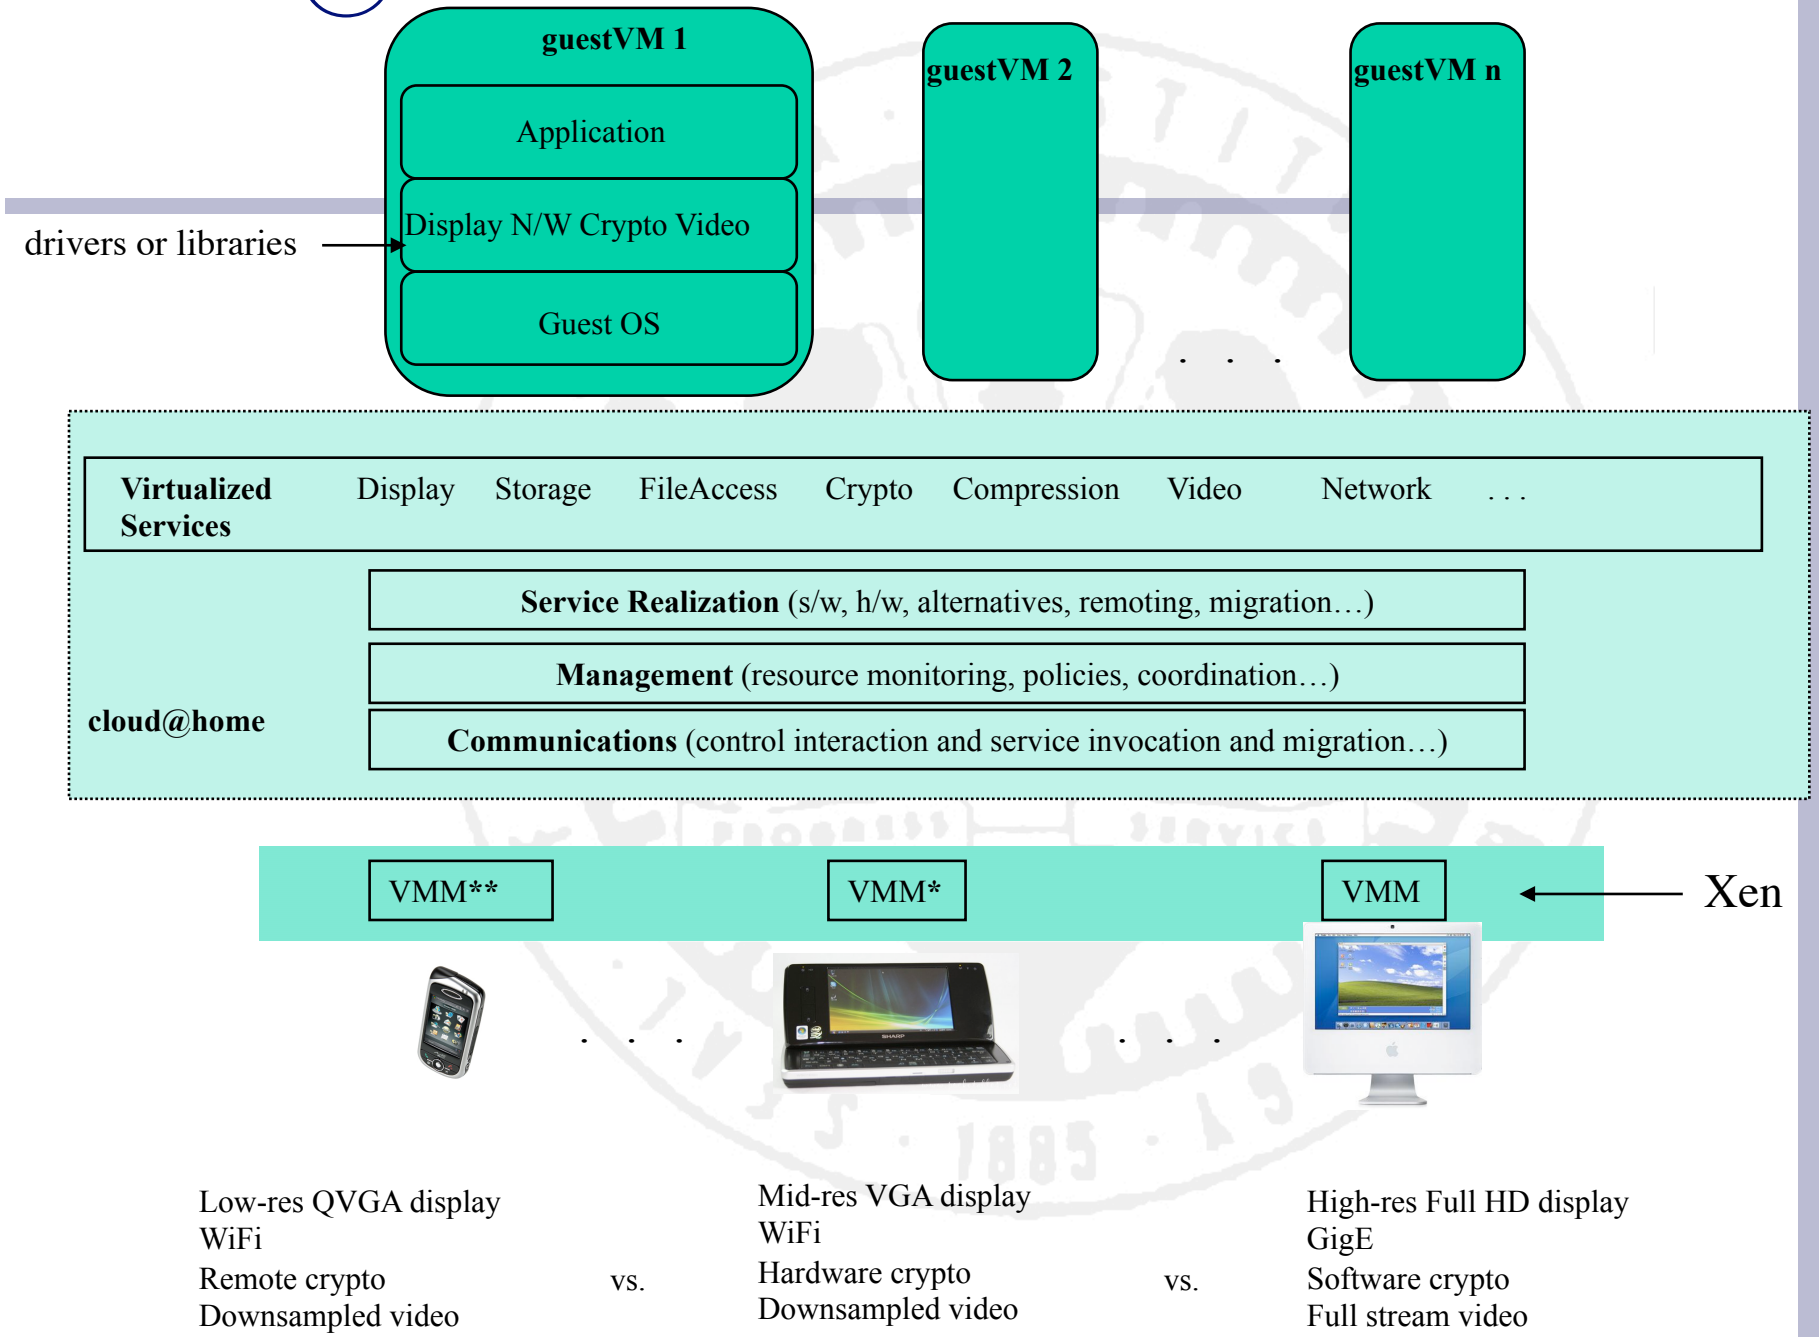
guestVM 1
Application
Display N/W Crypto Video
Guest OS
guestVM 2
guestVM n
drivers or libraries
Virtualized Services
Display
Storage
FileAccess
Crypto
Compression
Video
Network
. . .
cloud@home
Service Realization (s/w, h/w, alternatives, remoting, migration…)
Management (resource monitoring, policies, coordination…)
Communications (control interaction and service invocation and migration…)
VMM**
VMM*
VMM
Xen
Low-res QVGA display
WiFi
Remote crypto
Downsampled video
vs.
Mid-res VGA display
WiFi
Hardware crypto
Downsampled video
vs.
High-res Full HD display
GigE
Software crypto
Full stream video

Energy Policy

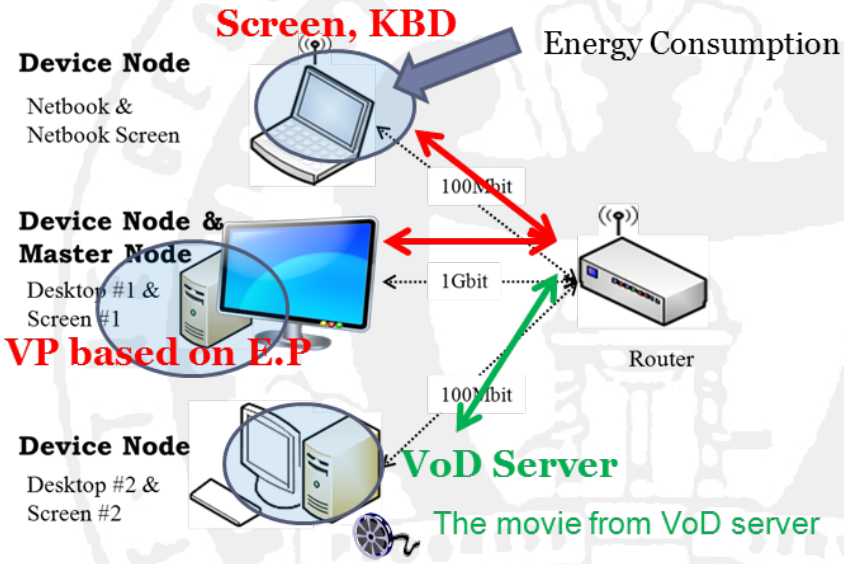

Configuration (Dom0+MVM)	Watt	Frame drop
1.66 GHz	14.3	21/1440
1.00 GHz	13.5	18/1436

Base line

Watt	Frame Drop
15.2	No
15.2	233/1441 (16%)

Saving 12.5%

Dealing with heterogeneous hardware

- Need to benefit from hardware accelerators, including SoC
- Virtualize at library level ala 'split-driver' model
- Virtualized Tolapai SoC
  - 'driver' and OCF backend in dom0
  - wrapper frontend library in guest VMs
  - extra latency in address translation and interrupt routing
- Additional opportunities
  - smart scheduling and resource sharing
  - remoting & emulation
- Other examples: GPU, GPS…