Architectures for Emerging Communication Processing Systems

Invited Talk - Oakland EastBay Chapter, IEEE Communications Society

September 21, 2006

Dilip Krishnaswamy (dilip@ieee.org), Ph.D.

Platform Architect, Intel Mobility Group

Overview

- Communication Processing Systems definition (for the talk)
- SoC platform architectures
  - Platform architecture optimization design vectors
  - 3D SoC platform architectures
- Platform architecture differentiation features
  - Security, manageability, virtualization
- Multi-protocol multi-channel multi-radio capabilities
- OverMesh: Network Centric Computing
  - Cross-layer Cross-overlay system architectures
  - Proactive Adaptive processing over heterogeneous networks

intel

Communication Processing Systems considered here

- System architecture for a node with integrated computing and communications processing

- System architecture with multiple nodes involved in distributed computation and communication over heterogeneous networks

intel

SoC Platform Architectures

- Optimize vectors such as performance, power, area, latencies, cost, energy/battery life, system efficiency
- Develop a high-performance low power embedded SoC arch
- Support multiple wireless protocols
- Take advantage of small protocol code size requirements
- Low latencies/high perf with integrated on-chip memory
- Consider options to keep memory closer to processors
- Choose memory size sufficient for comm and apps
- Dynamic power management with reduced leakage
- Energy-efficient performance architecture for comm and apps sub-systems

intel

Example SoC Platform: Intel® PXA800F Cellular Processor

Intel® 0.13µ Flash + Logic Process Technology

Intel® XScale™ Microarchitecture

Intel® Micro Signal Architecture

GSM/GPRS/EGPRS baseband communication logic

Integrated Flash and SRAM Memories -> 33% of chip area

Power Mgmt and Peripherals

Approx 10^8 transistors

12mm x 12mm TFBGA package

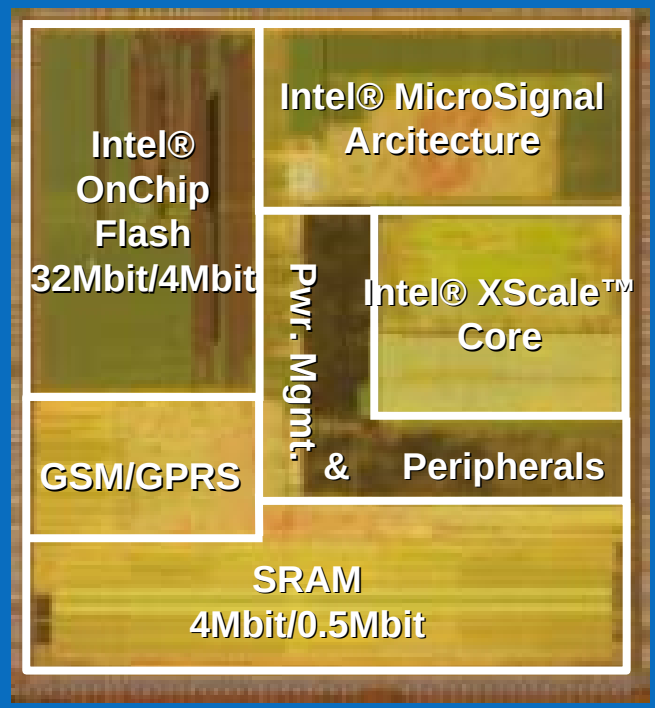

(Presented at IEEE HotChips 2003 and IEEE CICC 2003, Krishnaswamy et al)

Chip Modules

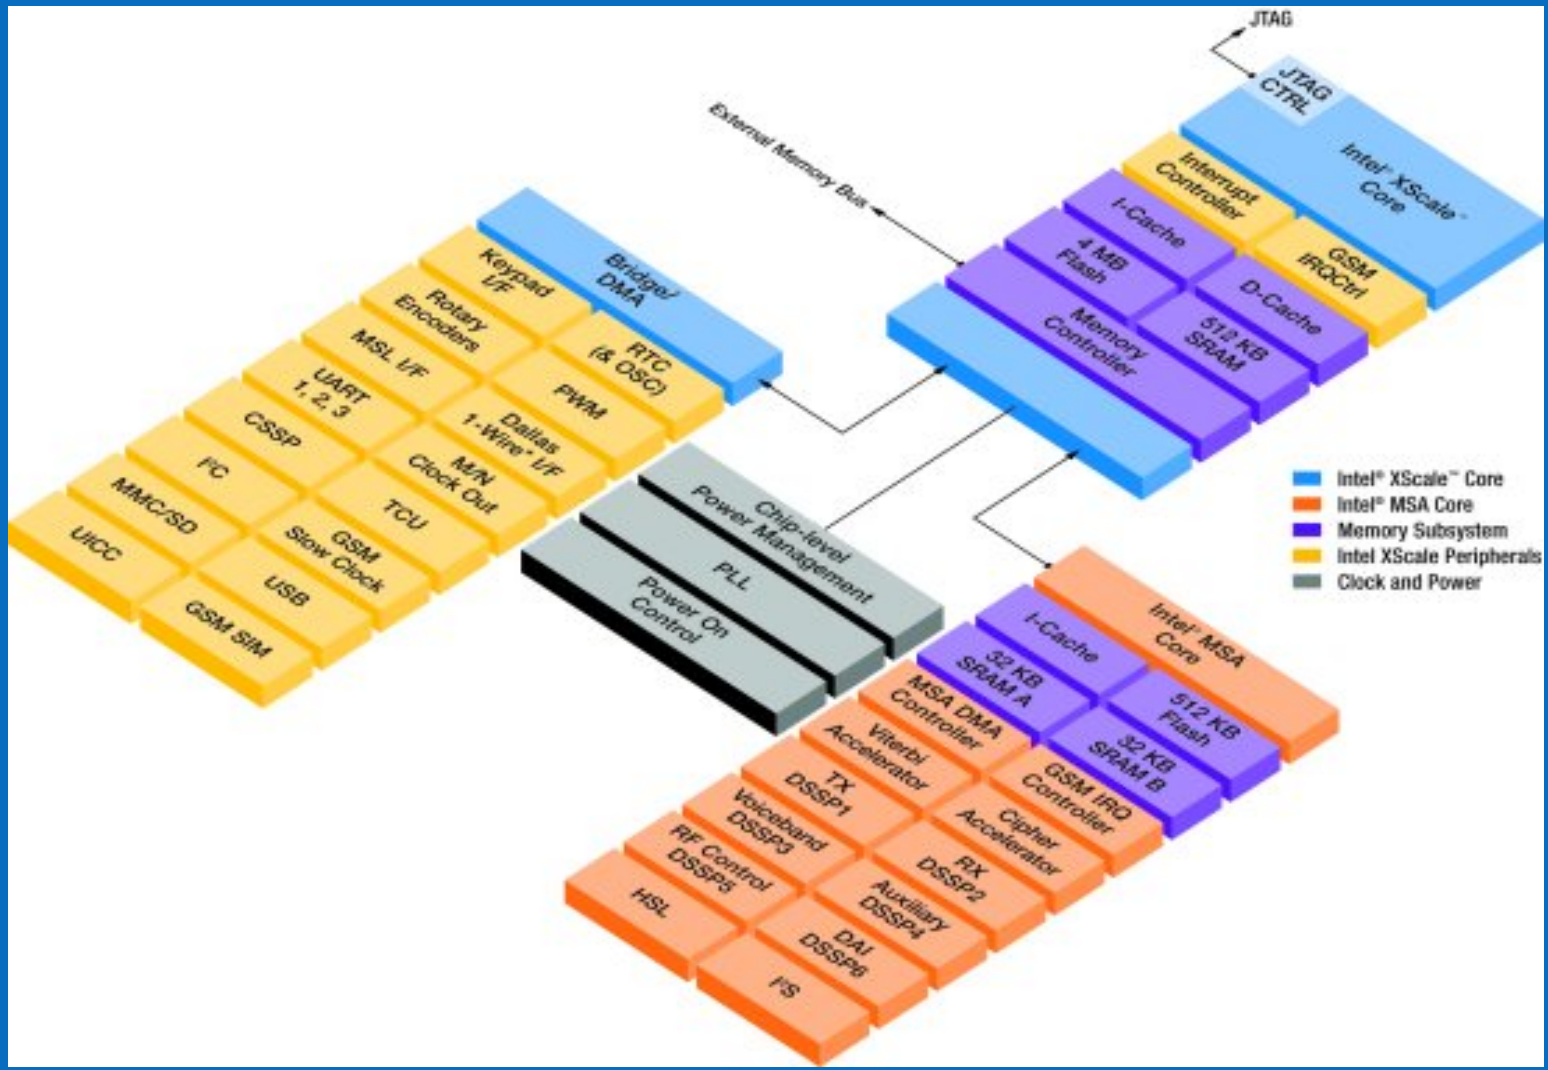

Intel® PXA800F/PXA800EF:
2.5G GSM/ GPRS Class 12/EGPRS Class 3

- **Complete Integrated baseband solution**
  - Intel® XScale™ Microarchitecture
    - ARM* V5TE compliant
  - Intel® Micro Signal Architecture
    - Modified Harvard Architecture
  - 36 Mb integrated Intel® OnChip Flash memory
  - 4.5 Mb integrated SRAM memory
- **Total solution**
  - GSM/GPRSClass 12/EGPRS Class3
  - L1-L3 Comm protocol software
  - RTOS Java-based platform
  - Designed for use with industry-proven mixed signal/RF solutions
  - Optimized PowerMgmt IC (Dialog)
  - Handset reference design (Elektrobit)

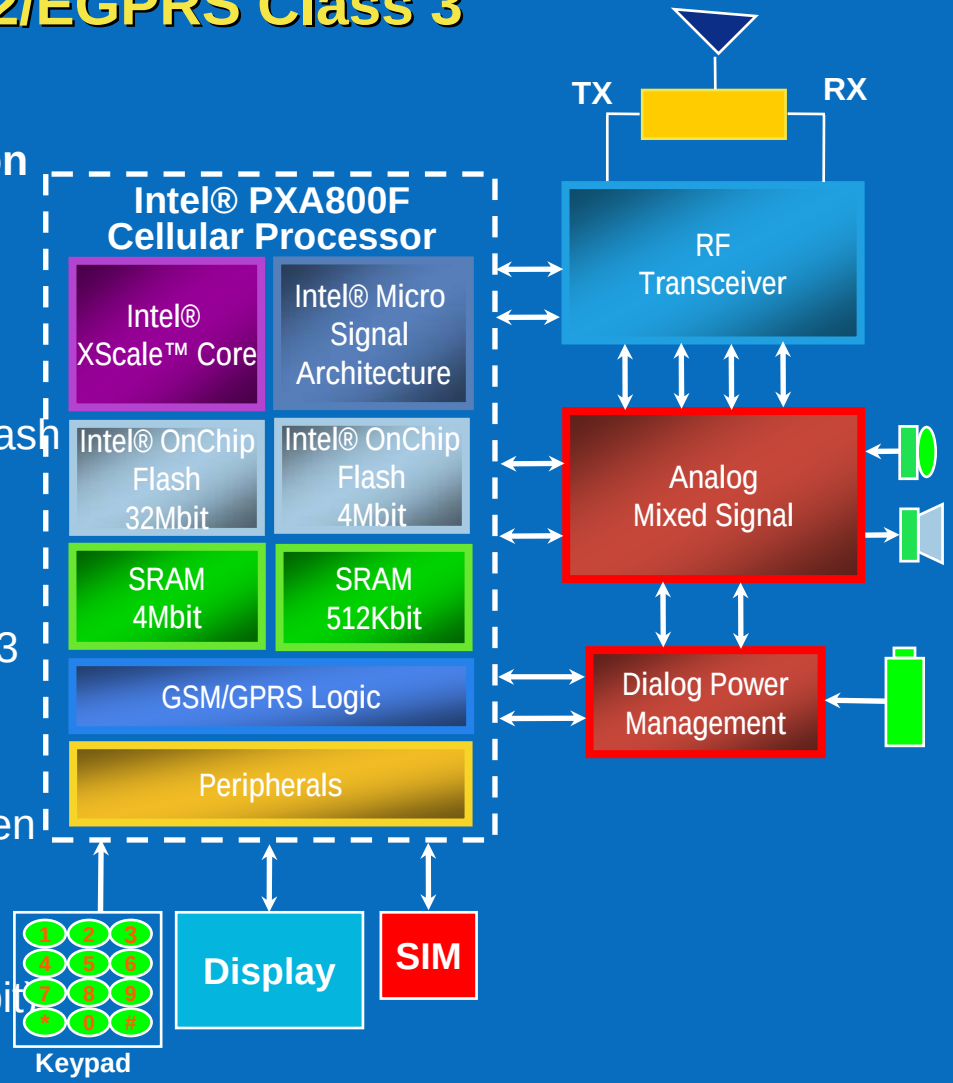

*Other names and brands are the property of their respective owners

Parallelism / Concurrency

- Wide independent buses from each core to integrated memories Internal memory buses provide direct accesses to memories; can be are wider and faster
  - Traditionally phone designs access memories from shared buses with arbitration penalties across narrow higher latency external memory interfaces which impacts performance
- Separate buses to access peripherals/ext mem

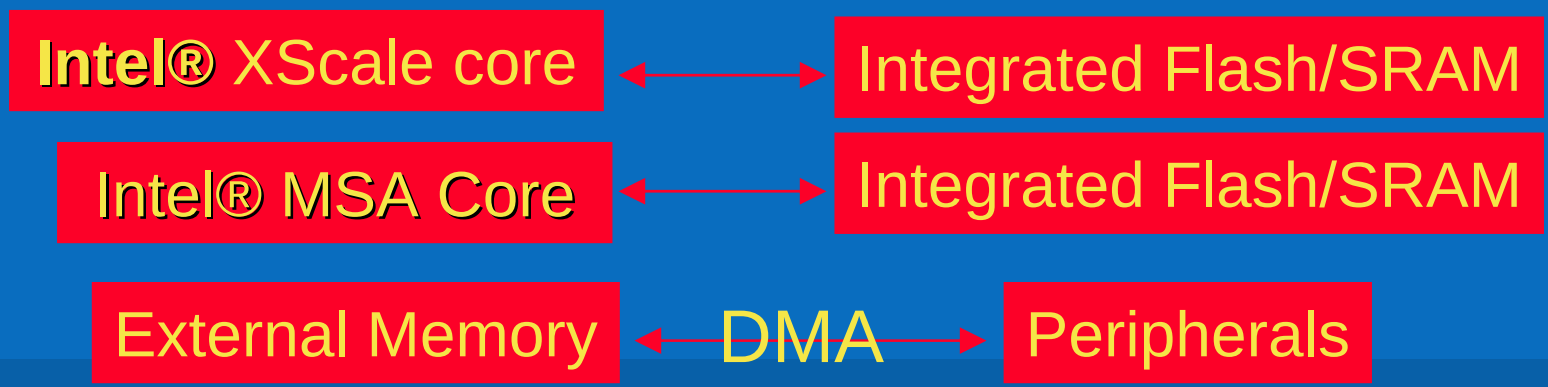

Reduced System Level Latencies

- System-level latencies to integrated memories greatly improved due to direct access and optimized design
- Critical-word first latencies in the Intel® XScale™ core sub-system as a %age of external memory access time

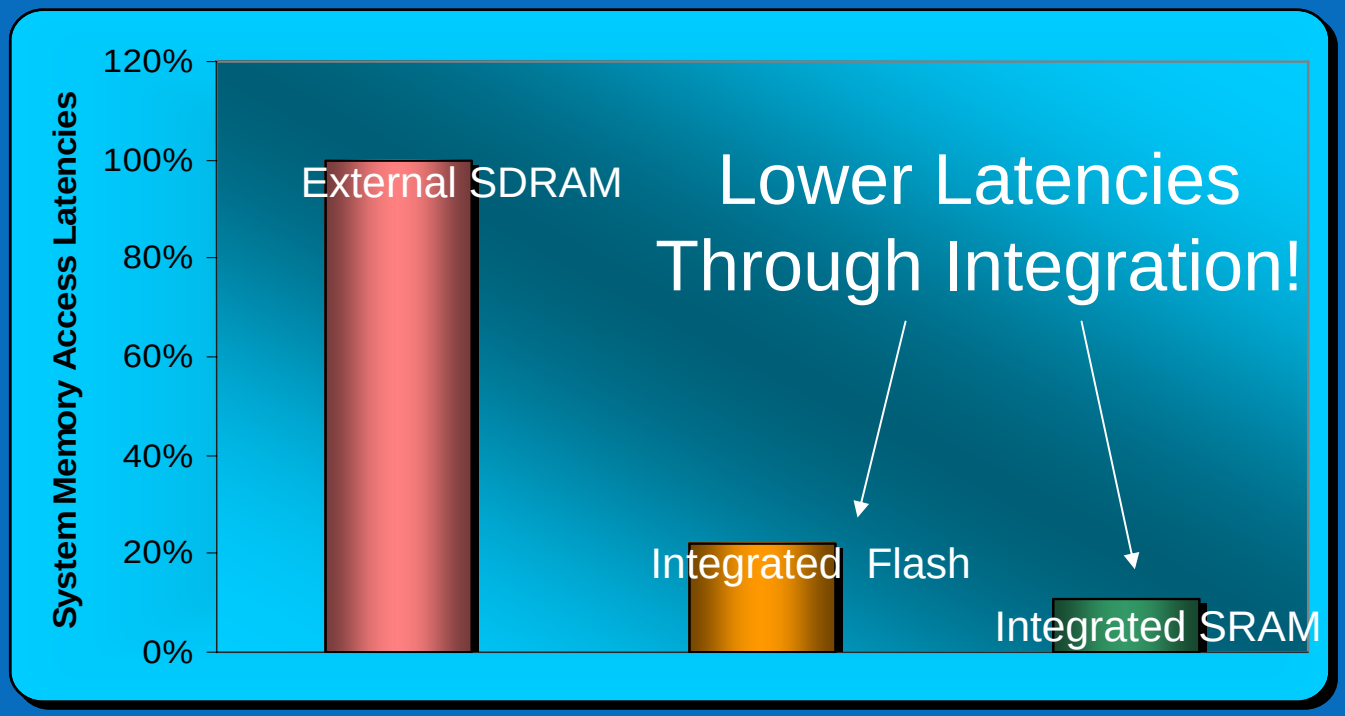

Active power usage optimization

- Total Active System level power is reduced when not using external memory
  - Power shown relative to system power when using external memory

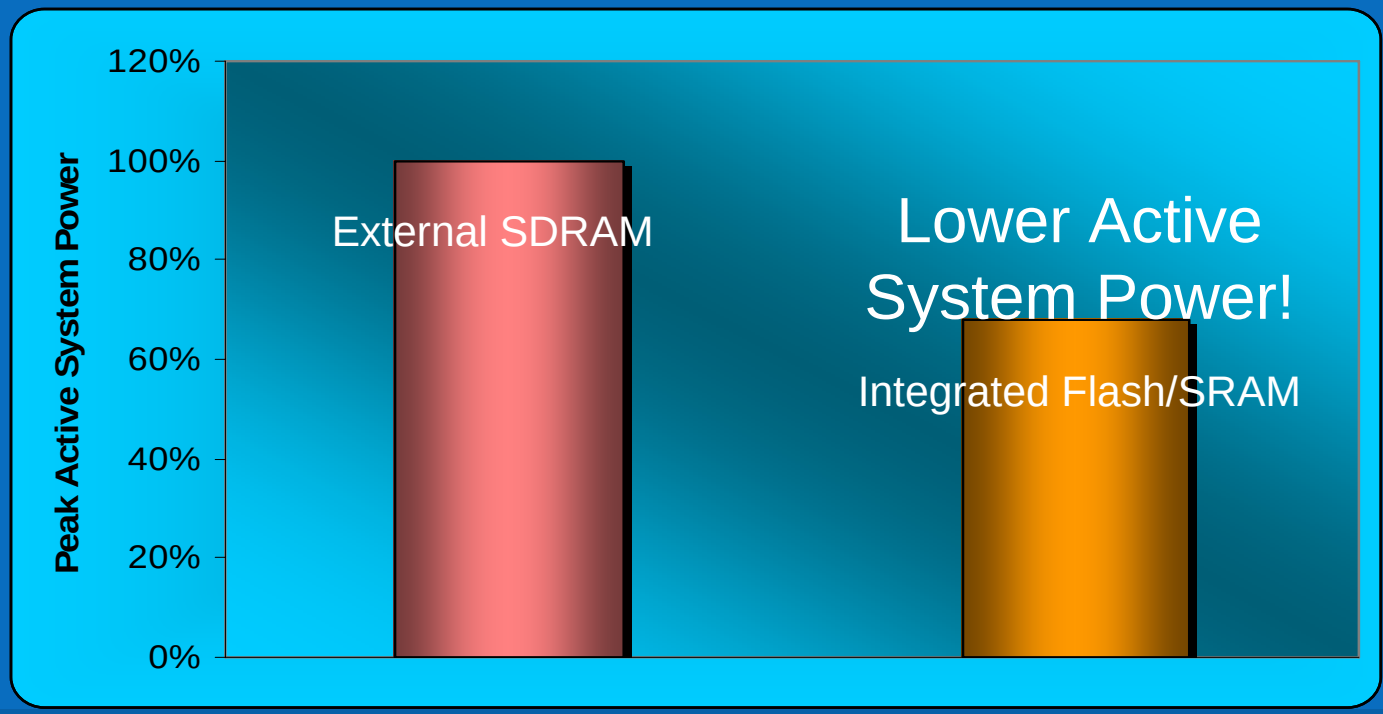

Energy Minimization

- Reduced Power * Reduced Latencies (Time)
  – ➔ Reduced Energy
- Lower energy needed for task completion
- System moves to a lower power mode earlier
- Power = VI (Leakage) & CV2F (activity) components
- Reduction in energy usage → increases battery life

intel

Power Modes

Mode	Cores	Peripherals	PLLs	I/O Ring	VCXO 13MHz Osc	RTC
Active	On	On	On	Active	On	On
Idle	Cores Inactive Clocks Disabled	Inactive (SW disable +Clocks disable)	On	Active	On	On
Standby	Cores Inactive Clocks Inactive Reduced Leakage States Retained	Inactive (except detection logic running with slow clock 32KHz)	Off	Inactive (sensing asynchronous inputs)	Off	On
Power Off	Powered Off (Real time clock active)	Off	Off	Off	Off	On

Energy-efficient power management

- Innovative power minimization strategies
- Active, partial-idle, idle, standby, and off power modes
- Internal power management unit to
  - Handle transitions between various power modes
  - Frequency scaling to reduce power
- Power mode to use determined by
  - Transition times/energy costs involved in transitions
  - Expected time in a lower power mode
- Hardware & software for optimal energy usage
  - Hardware state machines to handle transitions
  - Software decision-making to enforce transitions
- Custom power mgmt IC for additional power mgmt

intel

RF Edge-Rate Noise

- I/O digital buffers from the processor->memory interface exhibit a frequency spectrum determined by their edge-rates
- Faster edge-rates have broader frequency spectrum
- Spectrum can have RF components that interfere with the RF Rx channels
- Wireless data-centric phones suffer much worse due to extreme data transfer between a chip and external memories
- No I/O buffers in access path to integrated memories
  - → RF Edge Rate Noise Effects Reduced

intel

Benefits of memory integration

- Lower cache miss penalties
- Direct execution from integrated memories
- Higher System-level performance
- Lower power (less power wasted on the pins)
- Wider independent paths from memory to cache
- Outstanding response times during peak loading and task switching
- Greater Peak MIPS availability
- Reduces RF Edge Rate Noise Effects
- Reprogrammable Flash memory for communication and application sub-systems

intel

Performance Optimization

Worst case performance optimization needs to consider all possible data flows

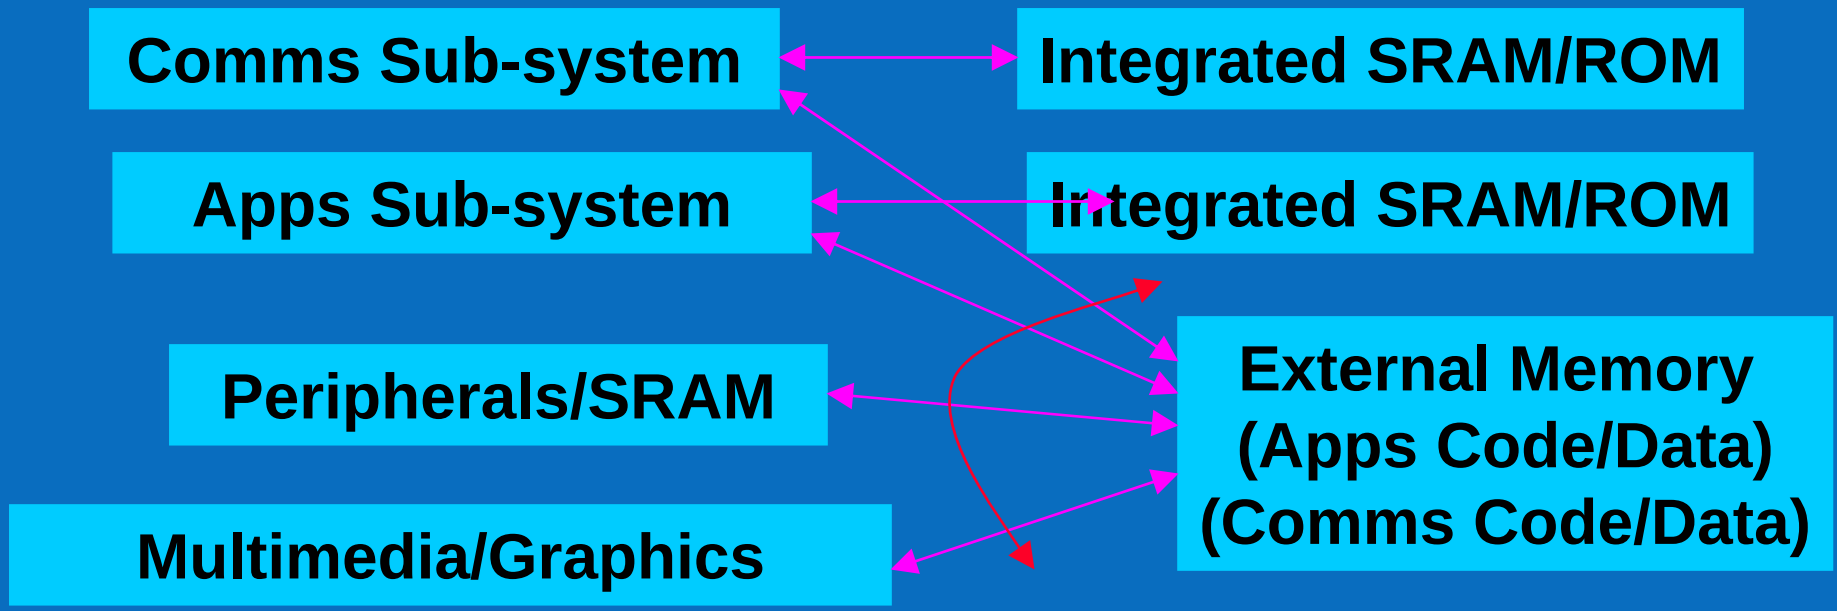

Key Architectural Considerations

- Parallelism
- Integrated memory size optimized for area/performance
- Reduced System-level latencies
- Higher System-level Performance
- Reduced System-level power
- Lower Energy
- Reduced RF Edge Rate Noise Effects
- Lower Cost for overall platform
- Reduced Area for the platform

intel

3D Architectures – the future

- Die-to-die interfaces, Through-Silicon-Vias (not merely conventional stacking)
- 3D extension of on-chip integration
- Reduced area footprint and more memory and logic possible
- Reduced distances, Parallelism, Energy/Performance Optimization

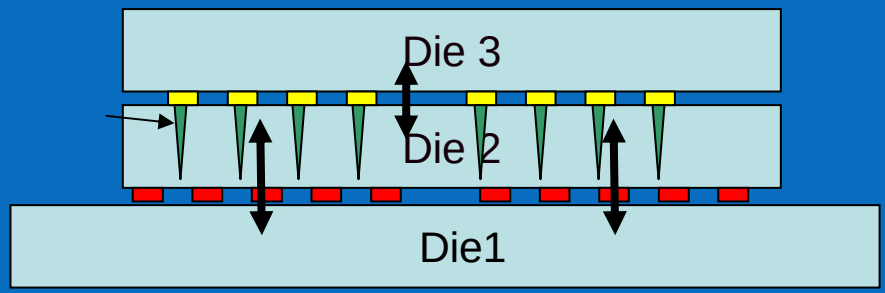

- See Ieong et al, IEEE CICC 2003, on 3D ICs, Tsui et al, IEEE Packaging Technology Conference 2003, on Thru-Silicon-Vias.

Secure Manageable Virtualized Platform Architectures

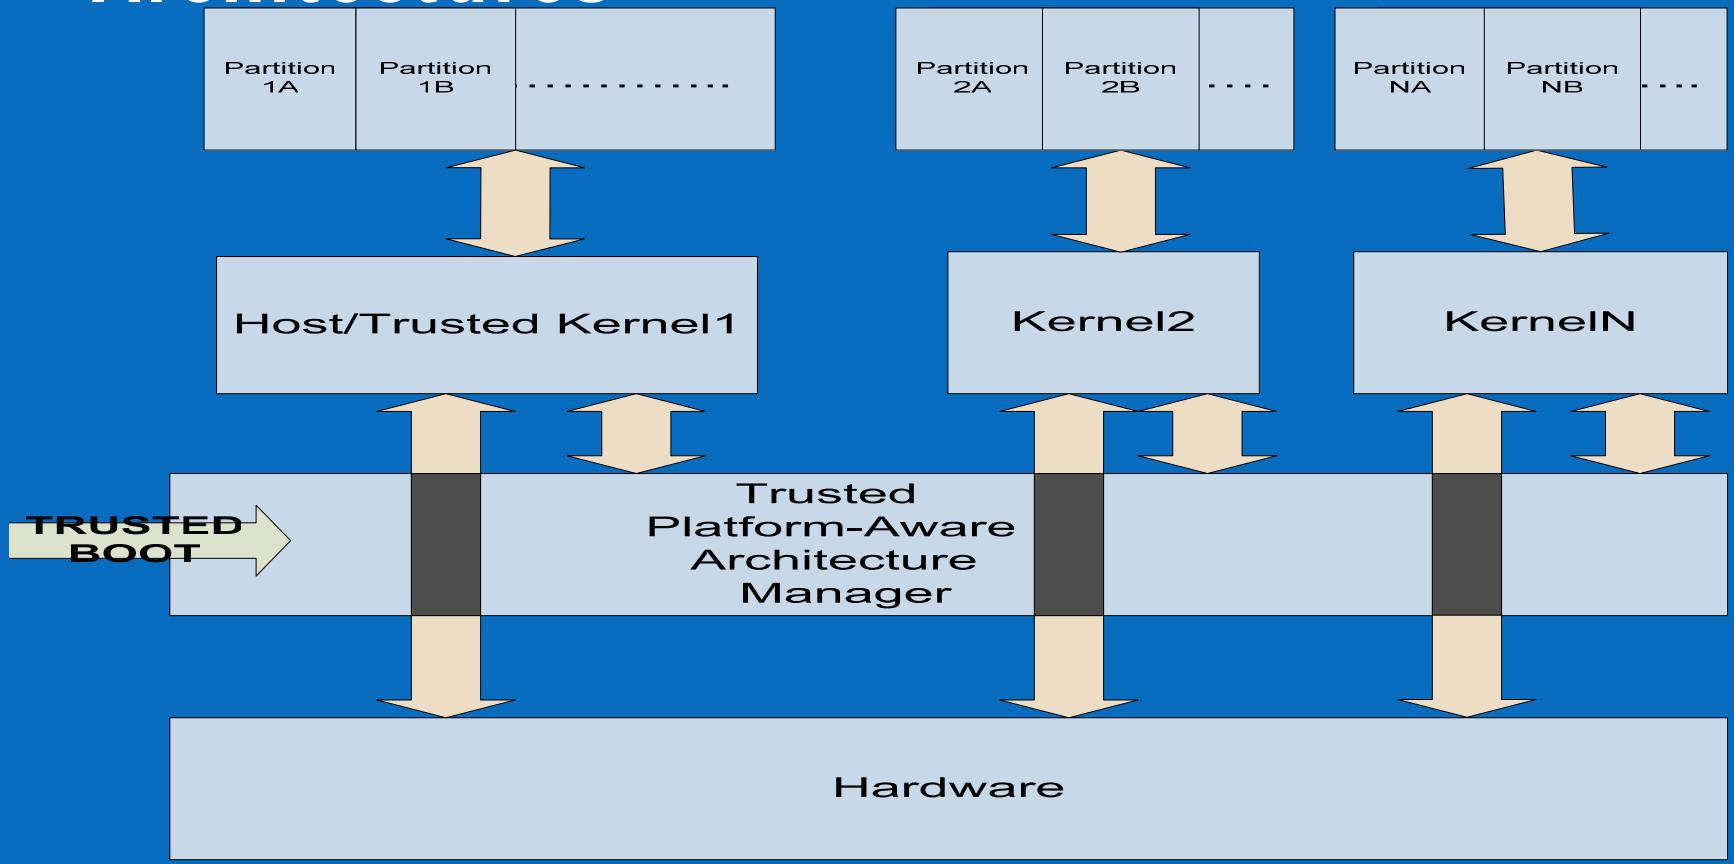

See IEEE ComSoc Magazine, Sep 2006, "Secure Manageable Mobile Platforms", Krishnaswamy, Hasbun, & Brizek

Useful Security Features

Transitive Root-of-Trust

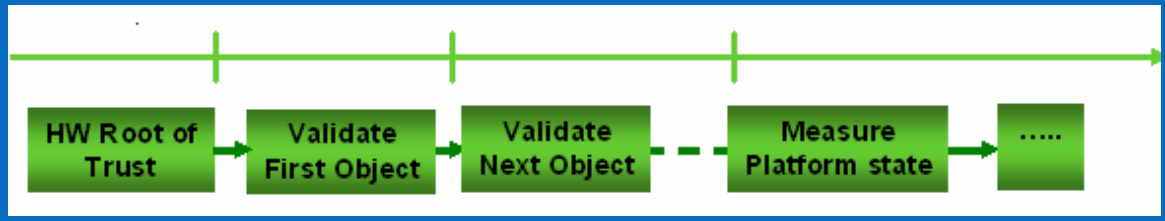

Key Hierarchies

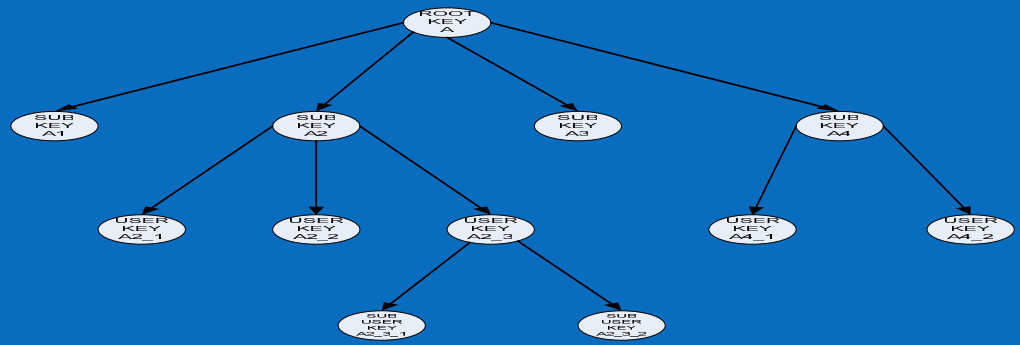

Key and Data-flow Management

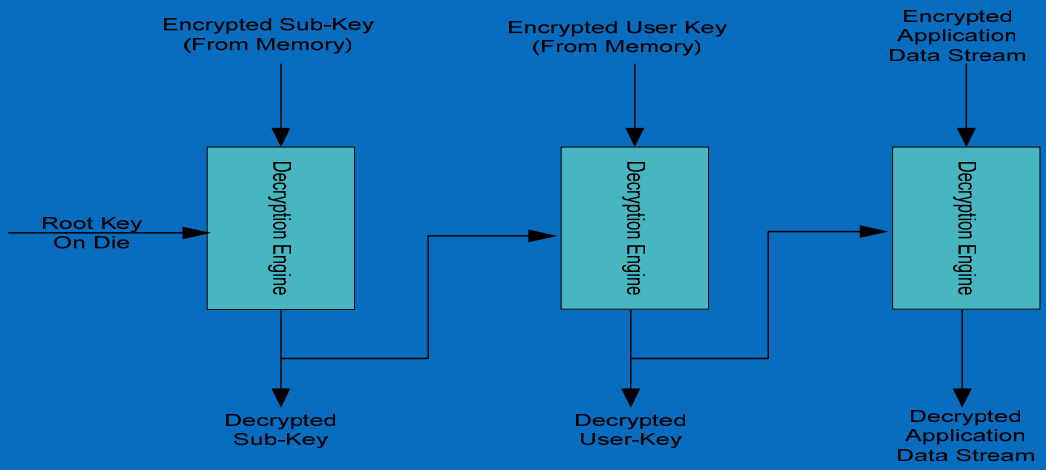

Wireless communications
Multiprotocol, Multichannel, Multiradio

- Multiple protocols
  - UWB for short range
  - WLAN for medium range
  - WiMAX/LTE/cellular for longer range
- Use multiple radios or reconfigurable radios for support for multiple protocols
- Use multiple channels with multiple radios for simultaneous communication

Multiprotocol Wireless communications

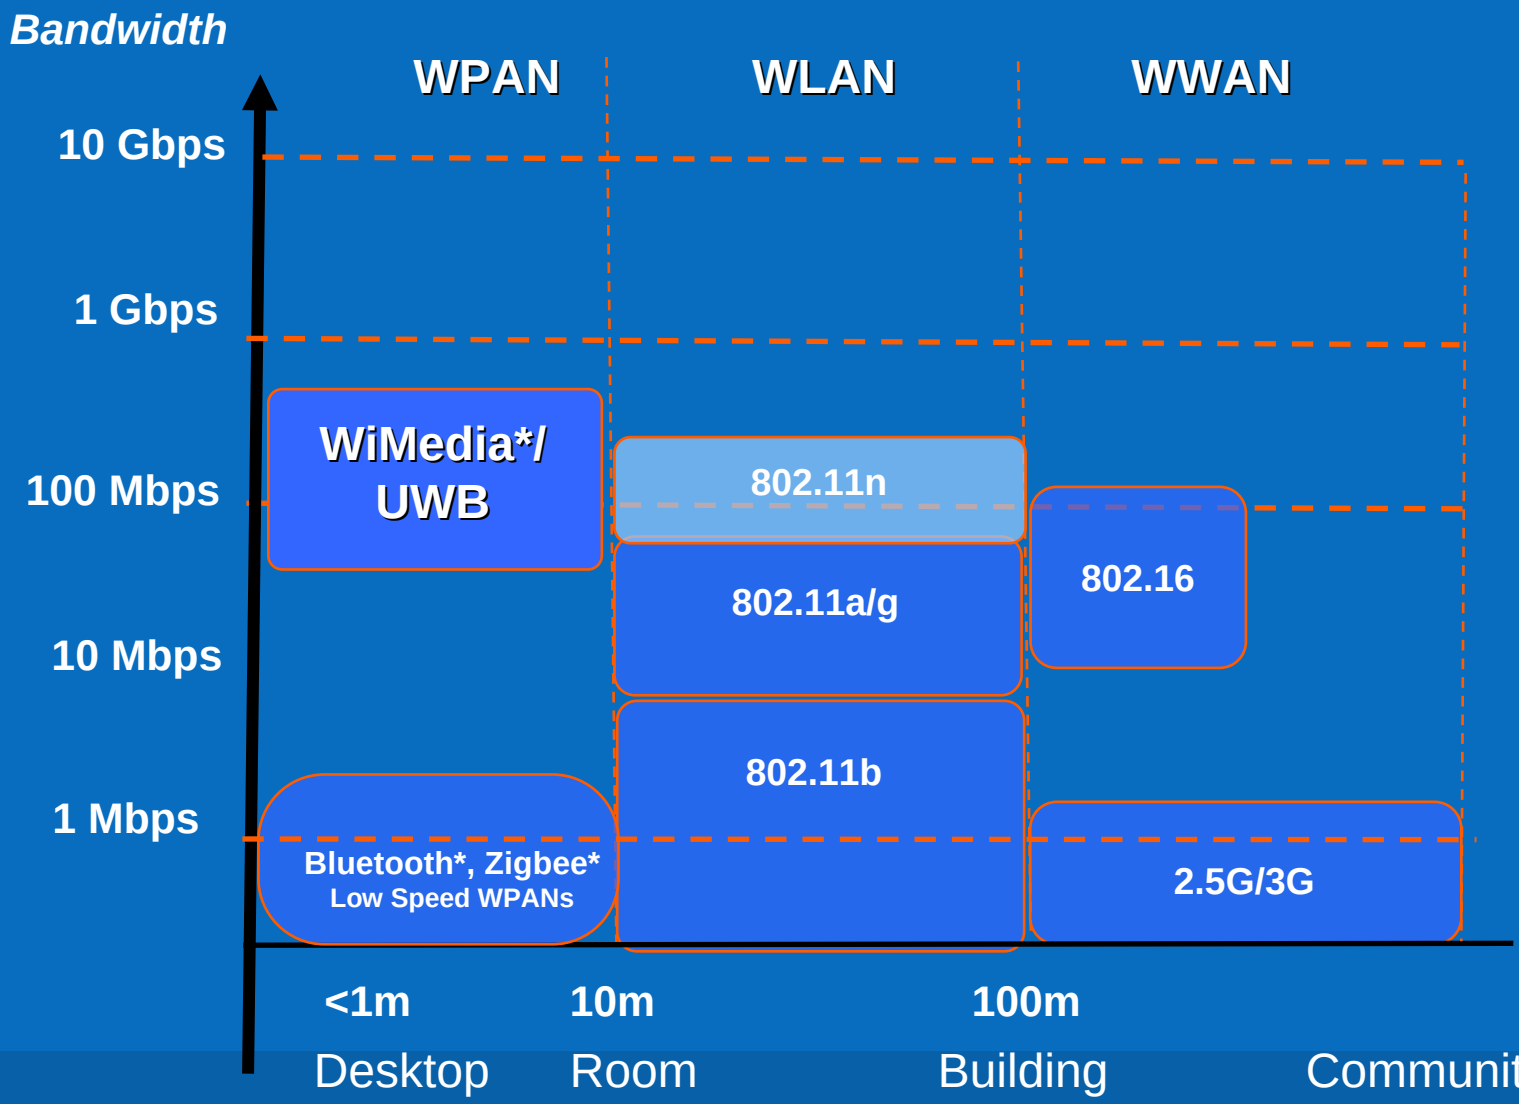

UWB – High-level view

What is UWB?

High-speed point-point wireless

Unprecedented 7.5 GHz span allocated

480Mbps @ 3m, 100Mbps @ 10m,
with multi-Gbps future

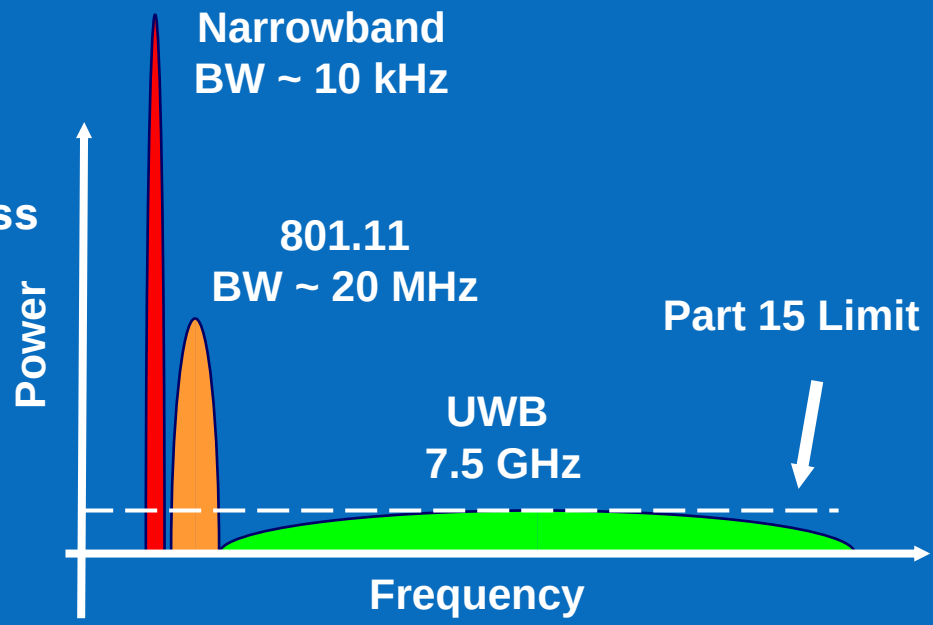

Why is it compelling?

Offers rich-content video streaming at low power

Complements and extends WiFi just as USB does for Ethernet

Can support Wireless USB, IP, 1394, and Bluetooth protocols

Can support up to 48 devices per host

802.11a/g WLAN protocol performance

- 75ns exponential fading matlab modeling for 802.11a/g PHY
- T(x)= A / (1 + e–λ(x– δ)) (Sigmoid modeling, Krishnaswamy, IEEE VTS 2002, 3G Wireless 2002)

54 Mbps = [(48 channels * 6 bits/symbol) / (4us/symbol)] * (3/4)

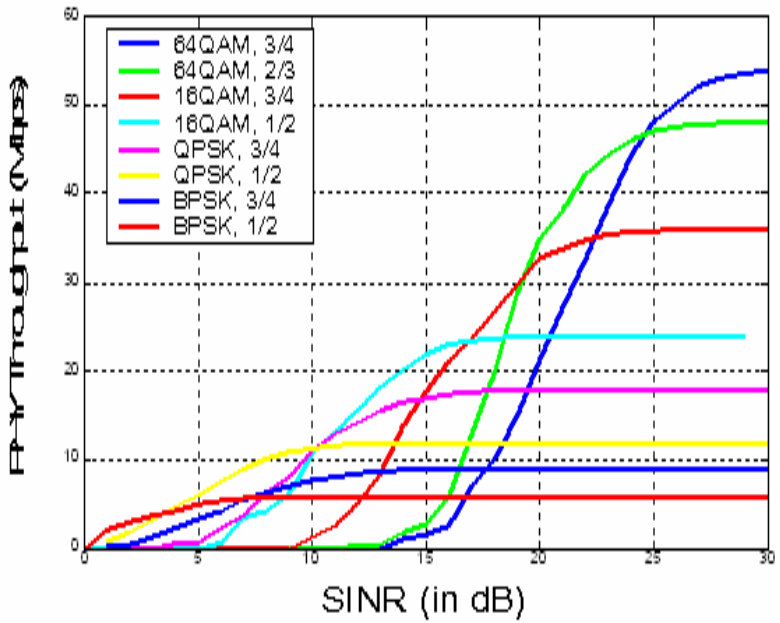

Figure 1: Sigmoid PHY throughput functions for 802.11a networks

Mod Scheme, code rate (k/n)	A (Mbps)	δ (dB)	λ (dB^{-1})	K_m (dB)
64-QAM, 3/4	54	21.2	0.419	29.45
64-QAM, 2/3	48	18.2	0.625	24.18
16-QAM, 3/4	36	15.1	0.352	23.27
16-QAM, 1/2	24	10.9	0.375	16.87
QPSK, 3/4	18	9.3	0.444	14.76
QPSK, 1/2	12	5.3	0.461	9.84
BPSK, 3/4	9	6.1	0.417	10.40
BPSK, 1/2	6	2.3	0.640	5.15

802.11n WLAN protocol performance

- Node optimizations at a single node may not be sufficient
- MIMO technology can deliver higher data rates
- Matlab modeling of a 2x3 MIMO channel Tx – SFC (Spatial Freq Comb), Rx – MMSE, Fading-75ns exponential

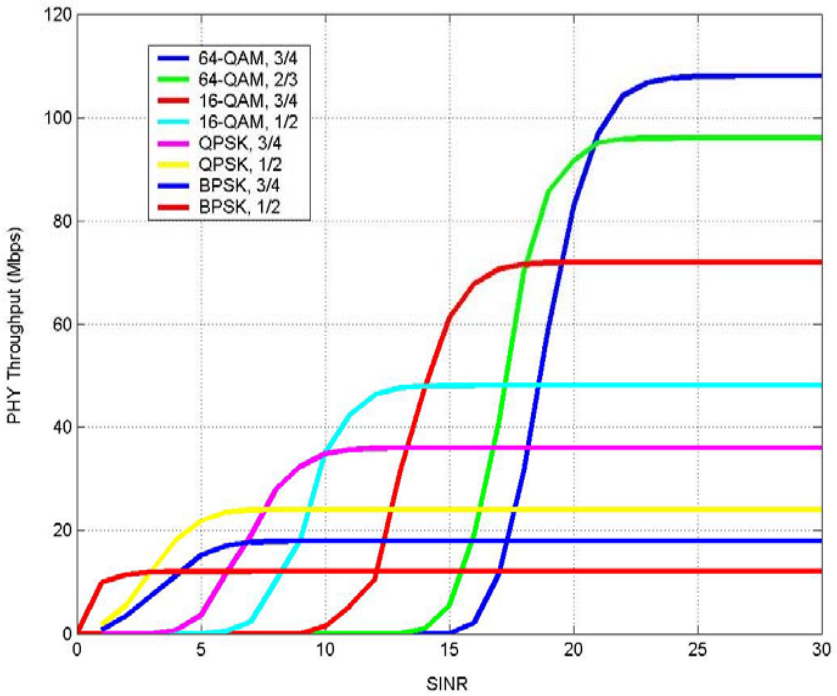

Mod Scheme, code rate (k/n)	A (Mbps)	δ (dB)	λ (dB^{-1})	K_m (dB)
64-QAM, 3/4	108	18.8	0.854	24.5
64-QAM, 2/3	96	17.4	0.833	23.1
16-QAM, 3/4	72	13.2	0.889	18.3
16-QAM, 1/2	48	9.4	1.258	12.9
QPSK, 3/4	36	7.1	0.907	11.3
QPSK, 1/2	24	3.1	1.006	6.6
BPSK, 3/4	18	3.5	0.815	7.2
BPSK, 1/2	12	0.7	3.158	2.0

Emerging 802.11n standard

- Current OFDM based systems
- 64pt FFT, 20 MHz channels (48 channels used for data)
- 2x3 MIMO systems → 108 Mbps
- Packet aggregation
- 40MHz channels -> 216Mbps
- 4x4 or 4x5 MIMO systems -> 432 Mbps
- 128pt FFT with 40MHz channels -> use more than 96 (48 *2) channels → 500Mbps
- Use QC-LDPC codes to approach Shannon limit → 72 * 8

  ->>> approaching 576Mbps

intel

WiMAX performance projections

Parameter	Downlink	Uplink	Downlink	Uplink
System Bandwidth	5 MHz		10 MHz	
FFT Size	512		1024	
Null Sub-Carriers	92	104	184	184
Pilot Sub-Carriers	60	136	120	280
Data Sub-Carriers	360	272	720	560
Sub-Channels	15	17	30	35
Symbol Period, T_S	102.9 microseconds			
Frame Duration	5 milliseconds			
OFDM Symbols/Frame	48			
Data OFDM Symbols	44			

Mod.	Code Rate	5 MHz Channel		10 MHz Channel	
		Downlink Rate, Mbps	Uplink Rate, Mbps	Downlink Rate, Mbps	Uplink Rate, Mbps
QPSK	1/2 CTC, 6x	0.53	0.38	1.06	0.78
	1/2 CTC, 4x	0.79	0.57	1.58	1.18
	1/2 CTC, 2x	1.58	1.14	3.17	2.35
	1/2 CTC, 1x	3.17	2.28	6.34	4.70
	3/4 CTC	4.75	3.43	9.50	7.06
16QAM	1/2 CTC	6.34	4.57	12.67	9.41
	3/4 CTC	9.50	6.85	19.01	14.11
64QAM	1/2 CTC	9.50	*6.85*	19.01	*14.11*
	2/3 CTC	12.67	*9.14*	25.34	*18.82*
	3/4 CTC	14.26	*10.28*	28.51	*21.17*
	5/6 CTC	15.84	*11.42*	31.68	*23.52*

Table 3: Mobile WiMAX PHY Data Rates with PUSC Sub-Channel[4]

[4] PHY Data Rate=(Data sub-carriers/Symbol period)*(information bits per symbol)

Copyright 2006 WiMAX Forum

intel

Future Wireless Networks

Vision - decentralized, scalable, and dynamically configured wireless network architectures

Scalability Opportunities:

- MultiTier Heterogeneous Mesh Networks, NodeInTheNetwork
- ClientAP, P2P Comm
- Multi-channel, Multi-radio, Multi-hop
- Reconfigurable multi-radios
- Distributed/Mobile/WiMAX portals

Challenges :

- Network scalability, Traffic Engineering
- End-to-End QoS (Bandwidth/latency)
- Load balancing with reduced congestion at portals (wireless/wired interfaces)
- Flexibility, Adaptivity, Fault tolerance

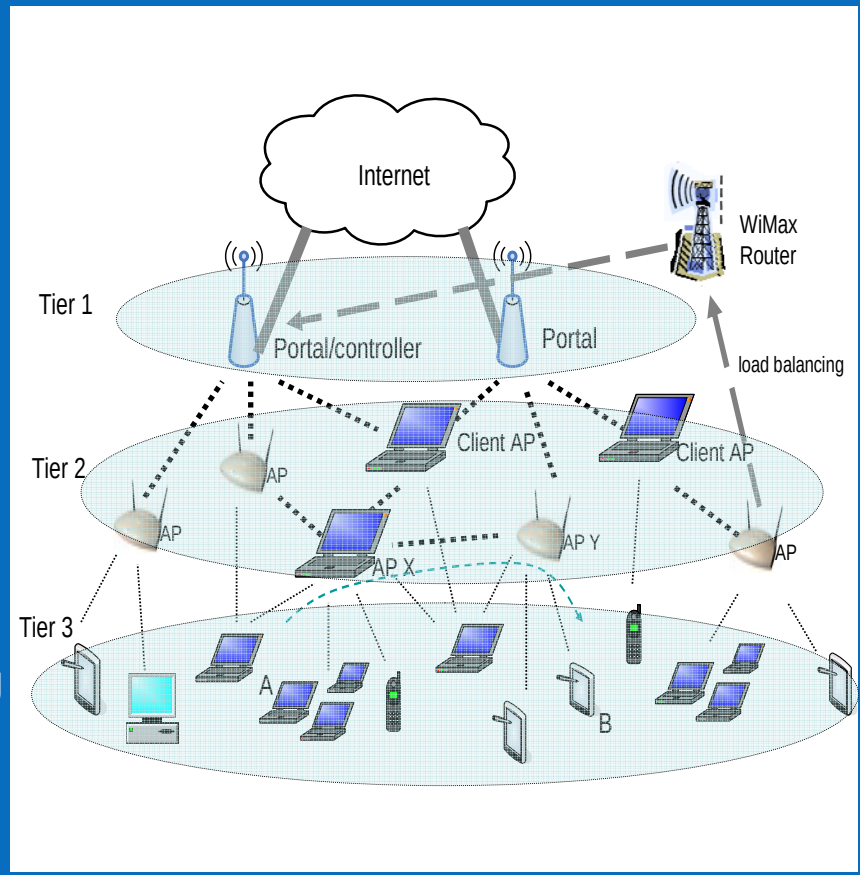

Network Infrastructure Today:

Centralized, Fixed Network Infrastructure

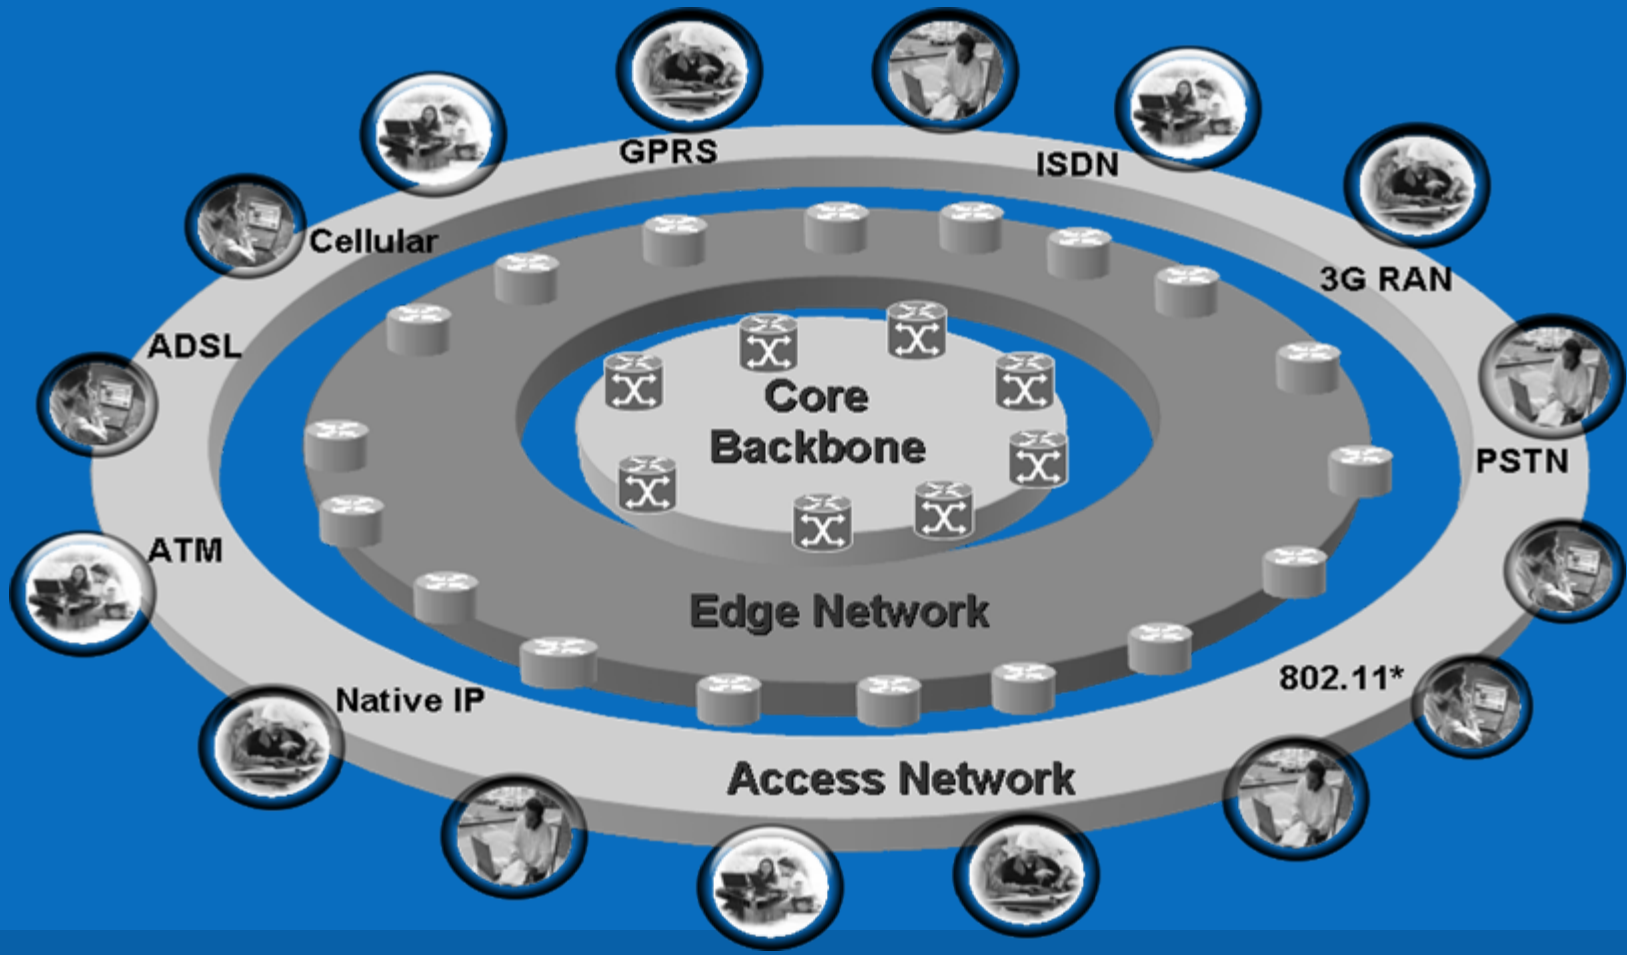

OverMesh - Network Centric Computing

A decentralized architecture of communications with a highly virtualized, converged computing and communications node-based wireless mesh architecture, having highly intelligent management capabilities.

User embedded inside the network will be able to freely associate services and to leverage resources such as computing, storage, and bandwidth to rapidly advance network innovation.

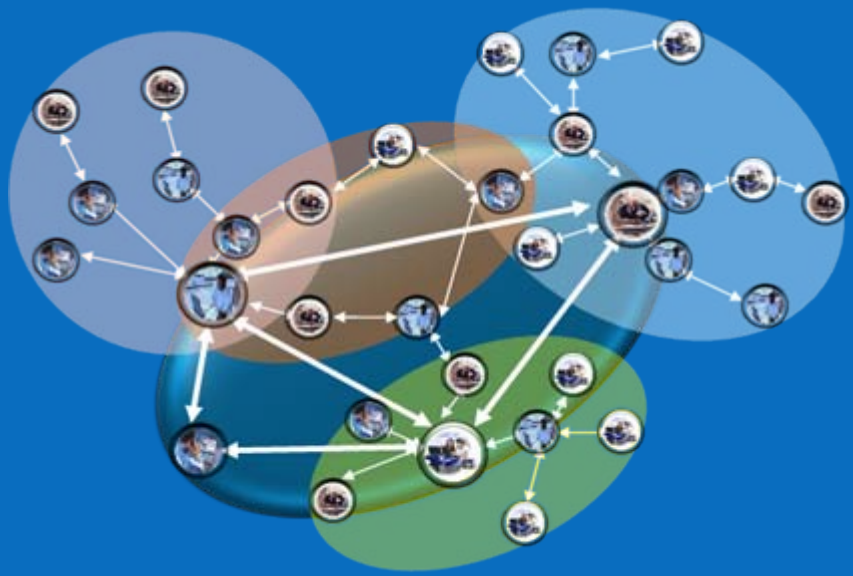

Sparse OverMesh networks

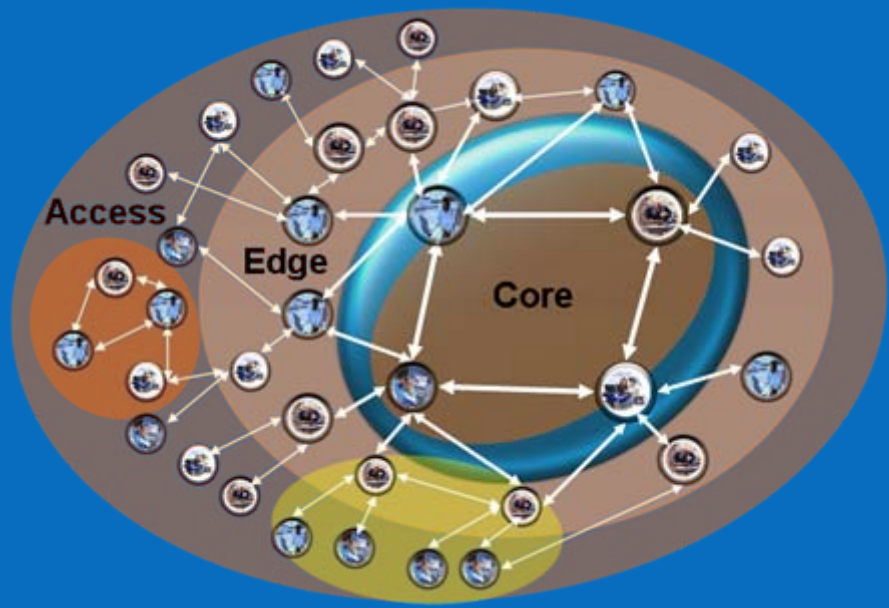

Structured OverMesh networks

See, OverMesh: Network Centric Computing, To appear IEEE Communications Magazine, J. Vicente, G. Ding, S. Rungta, D. Krishnaswamy, W. Chan, K. Miao

intel

Proactive Adaptive Cross-layer Cross-Overlay Architectures

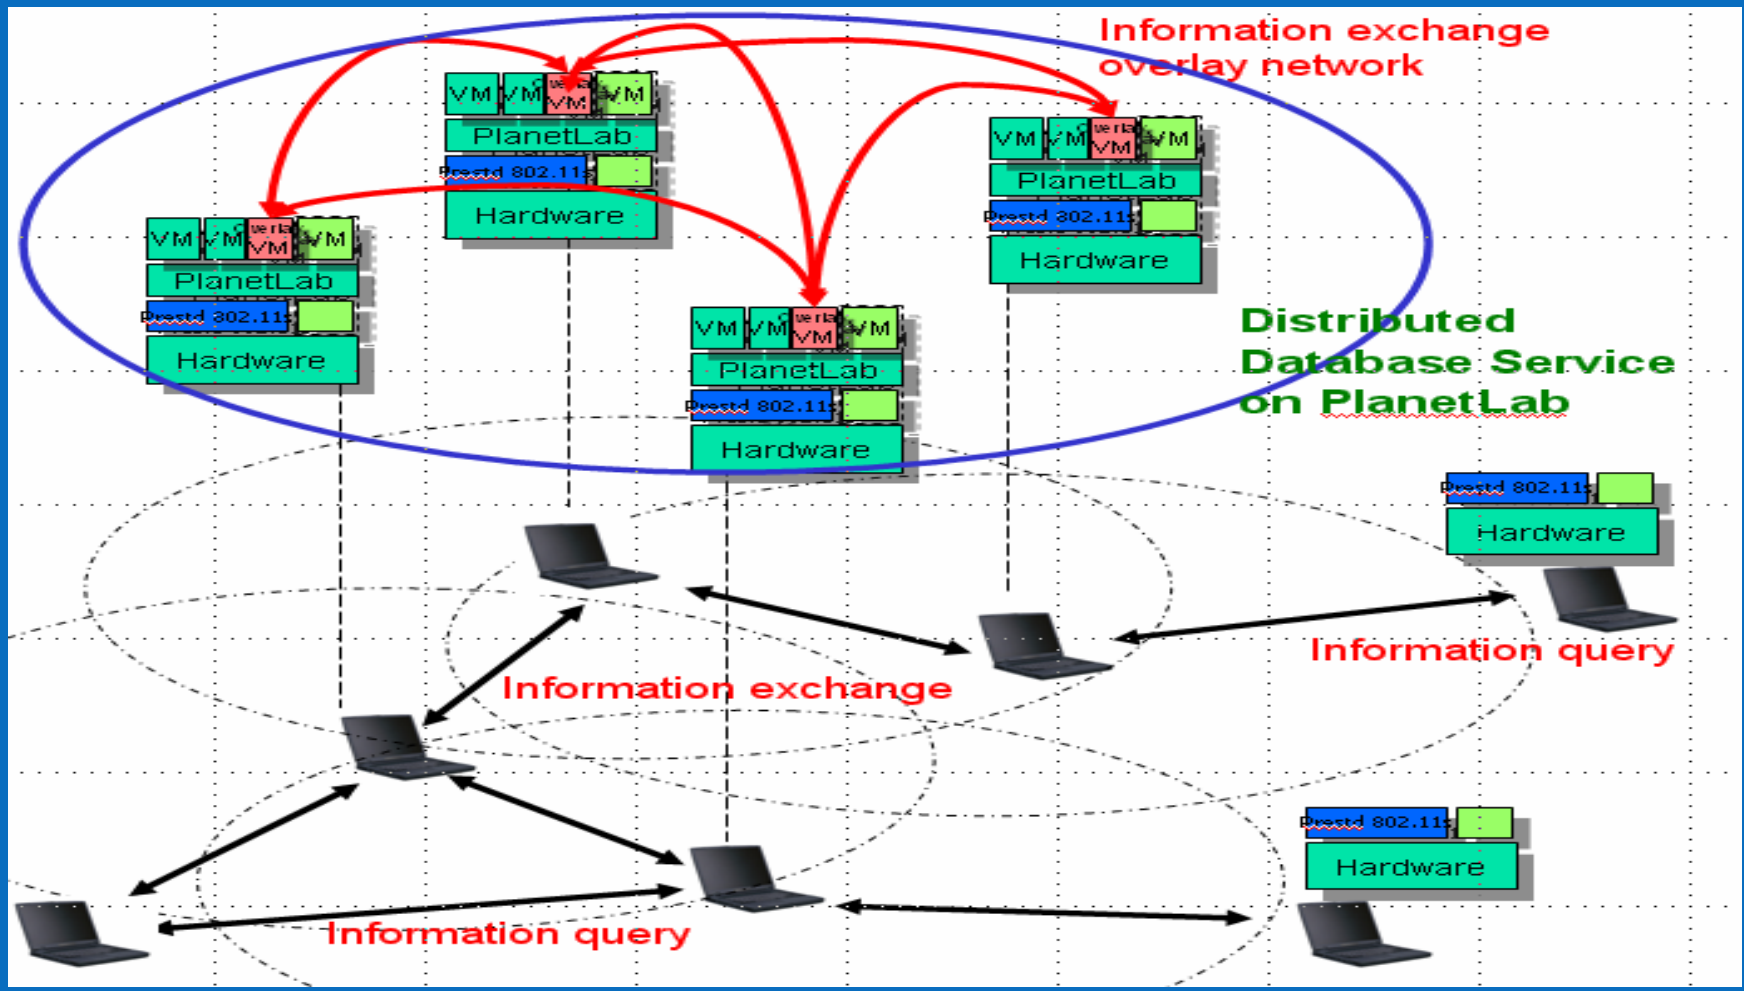

(recent research see IEEE WiMesh Workshop, Sep 2006, Krishnaswamy et al)

Node-level System Architecture

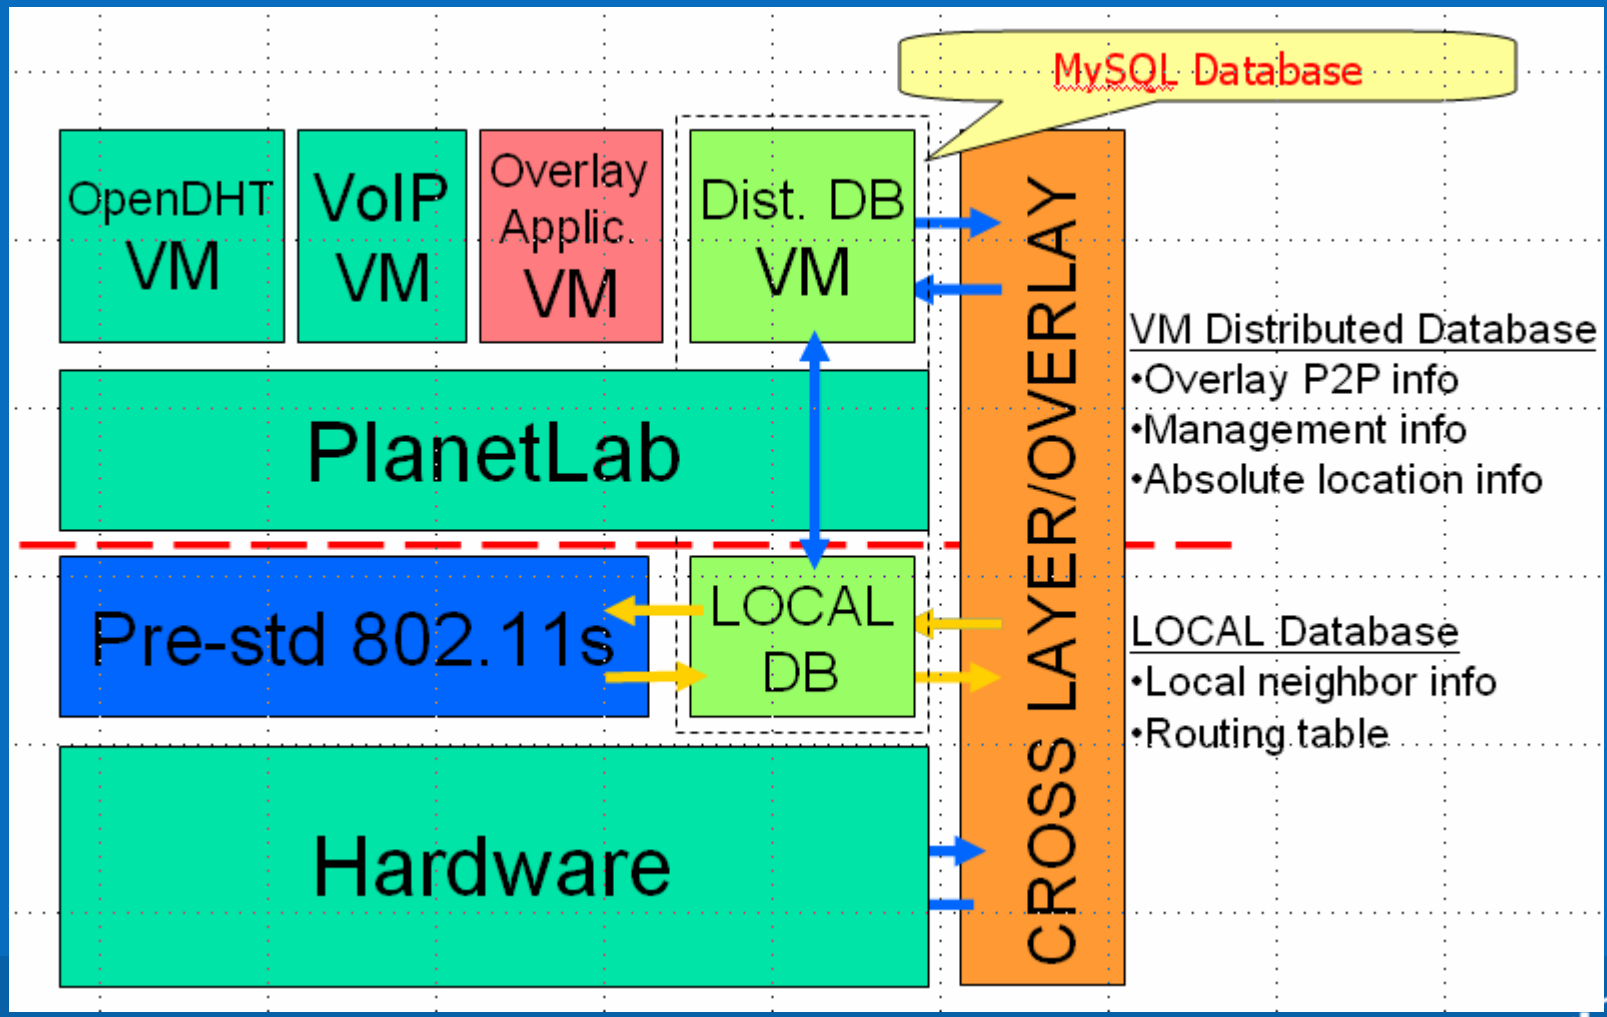

Network-aware proactive optimizations

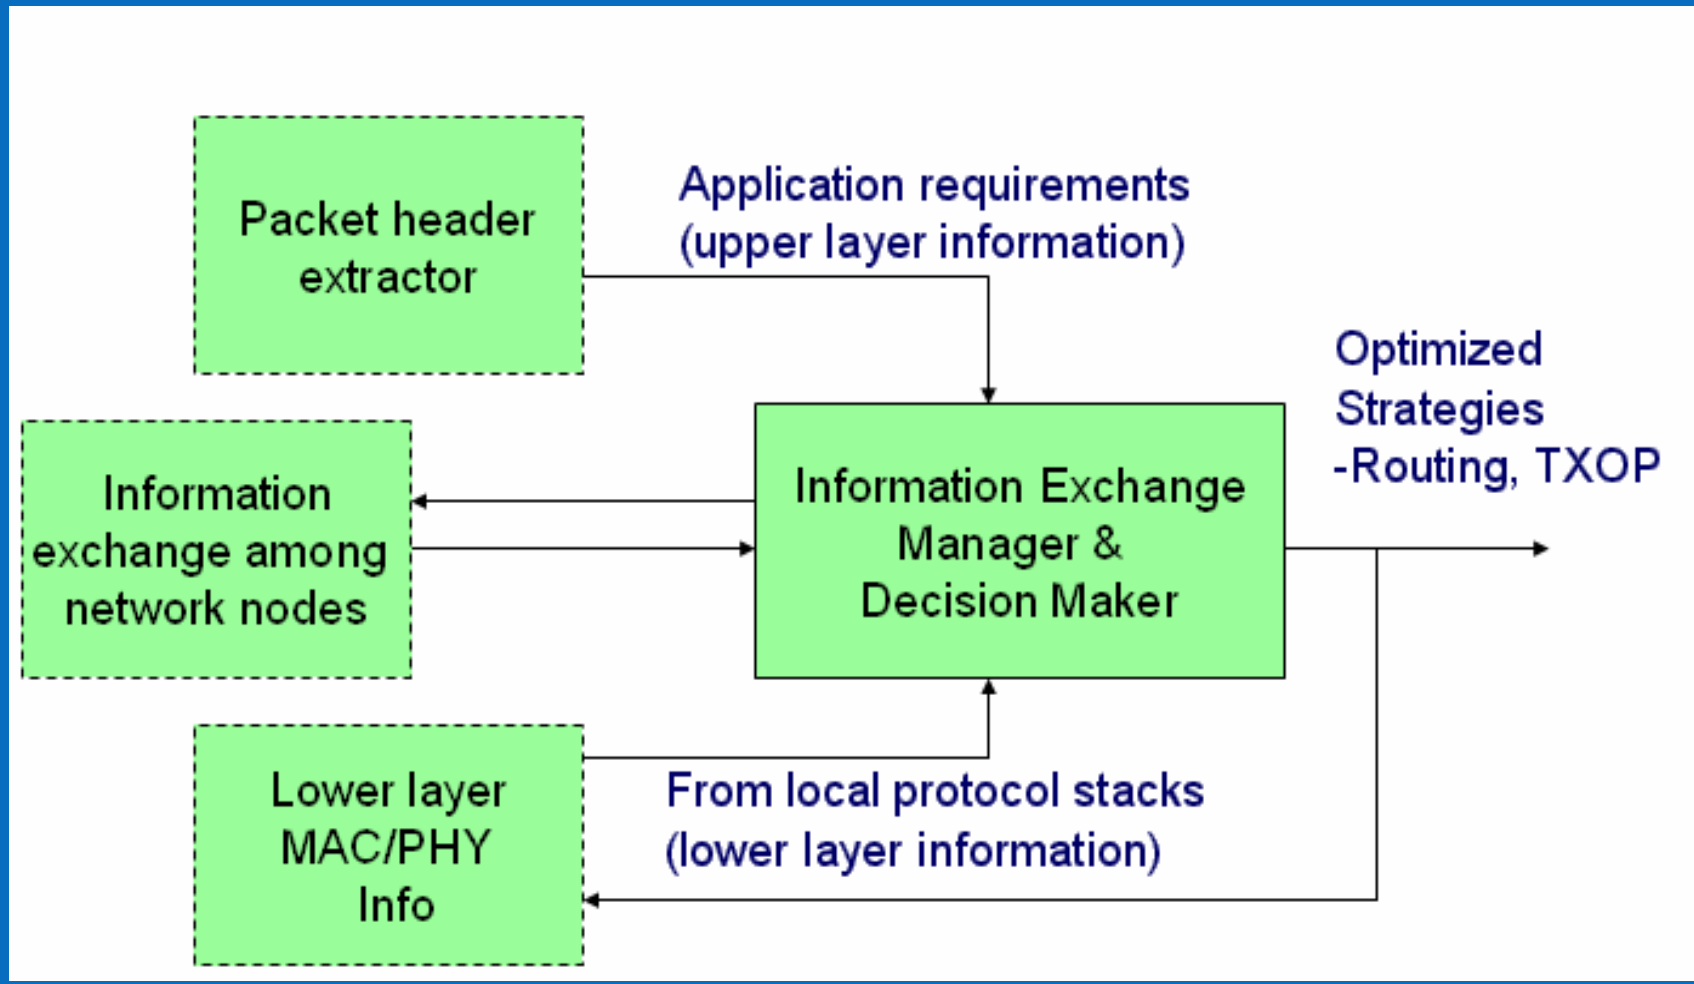

Distributed Network Information Base Design

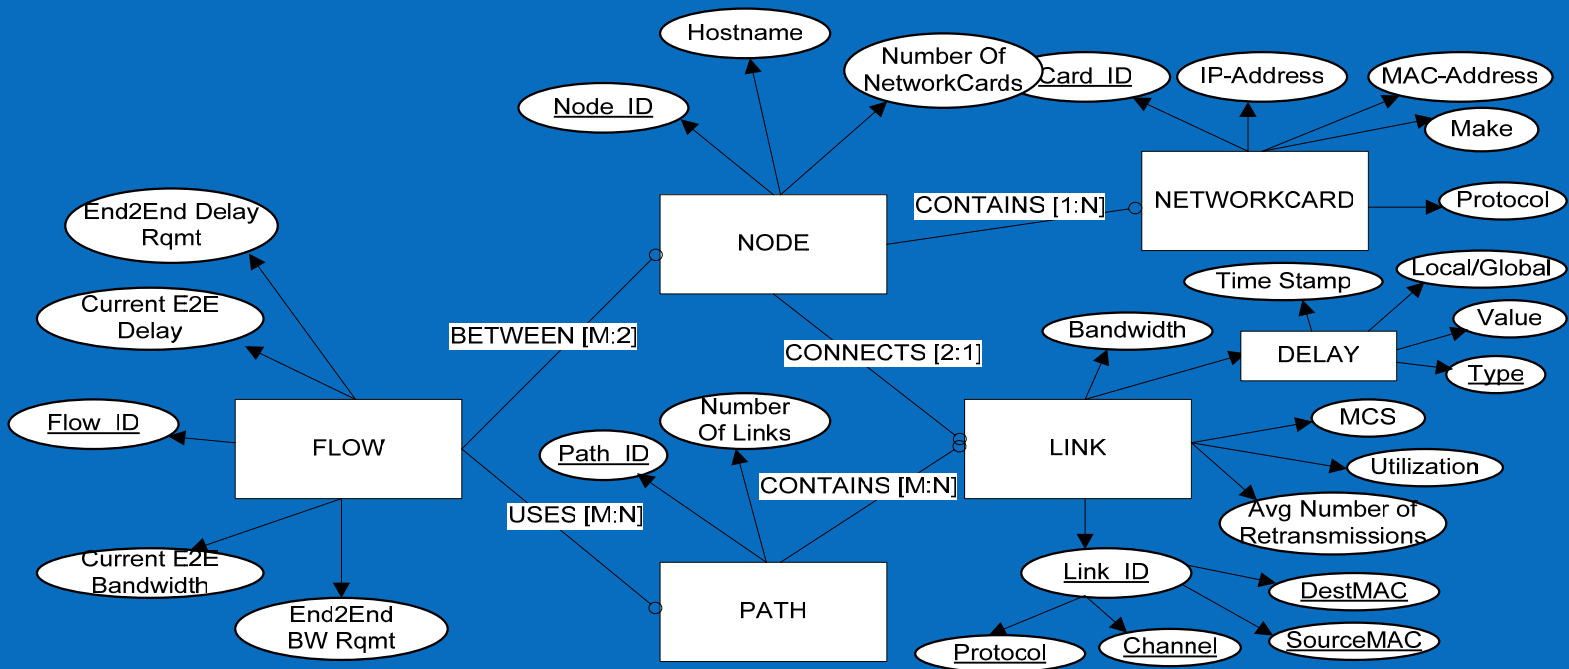

Application – Proactive Multi-hop L2 routing

$$\underset{r \in RouteGraph}{\text{minimize}} \sum_{k} \left(1 - \frac{TH_{k,r}^{overall}}{demand_k} \right)^2$$

$$\text{subject to} \sum_{i \in path_k^n} ETT_{k,i} \leq deadline_k - gap_k$$

$$\sum_{k} t_{k,i} \leq 1$$

$$TH_{k,r}^{overall} \geq demand_k, \quad \text{if } f_k \text{ is a nonscalable application}$$

$$path_k^n \in r, \quad \text{for all } f_k$$

intel

Mesh routing example

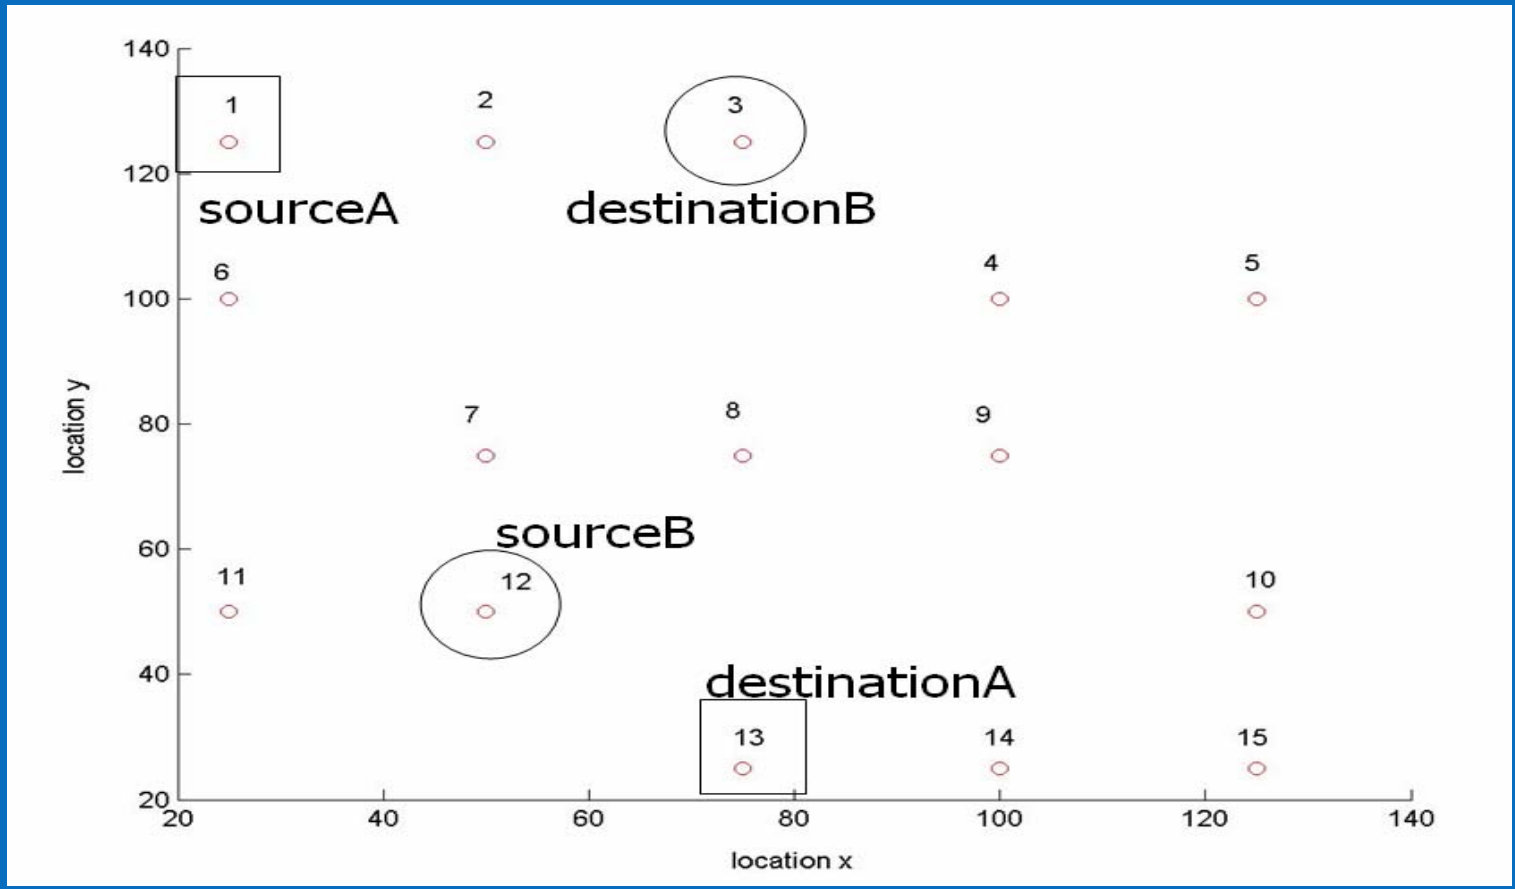

Mesh routing example contd.
- multi-path routing

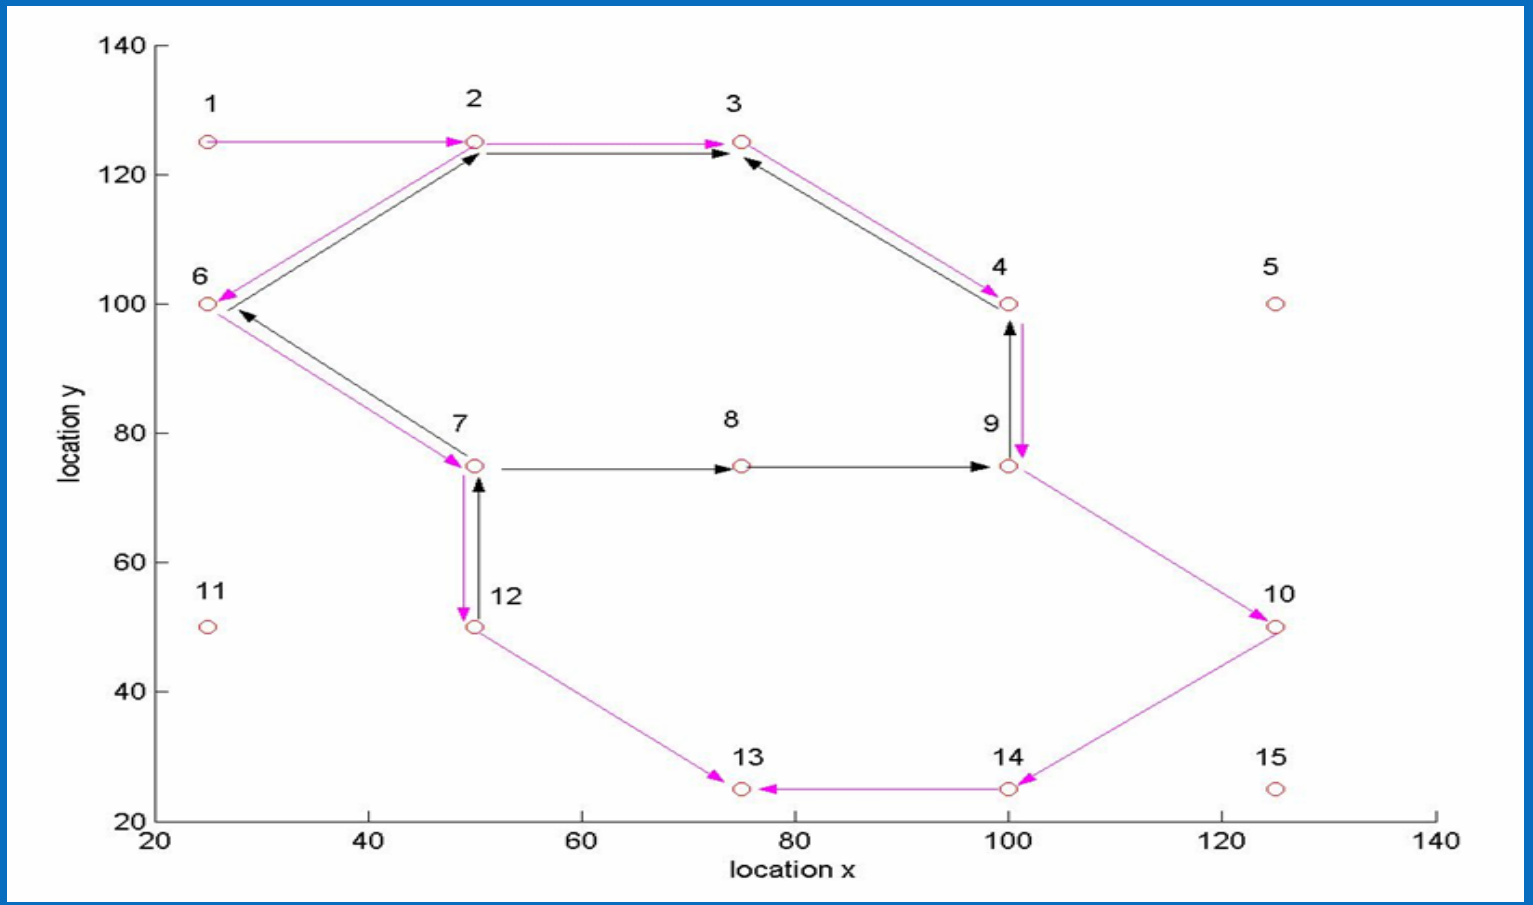

Simulation results – good channel PER (0.05 to 0.15)

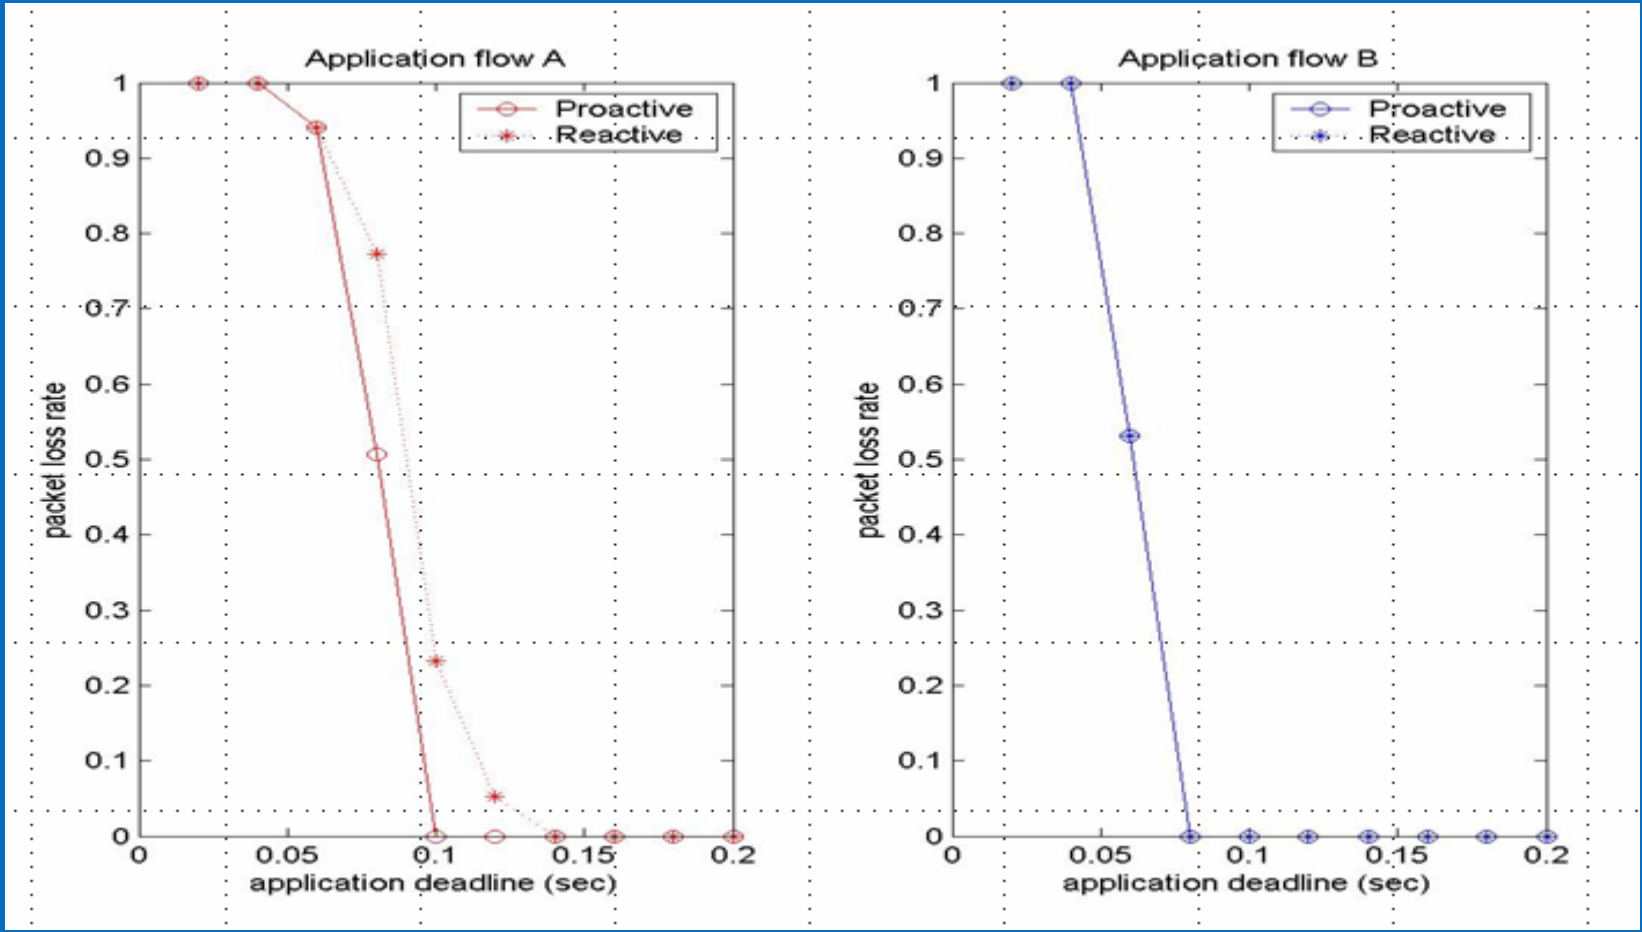

Simulation results – worse channel PER (0.05 to 0.30)

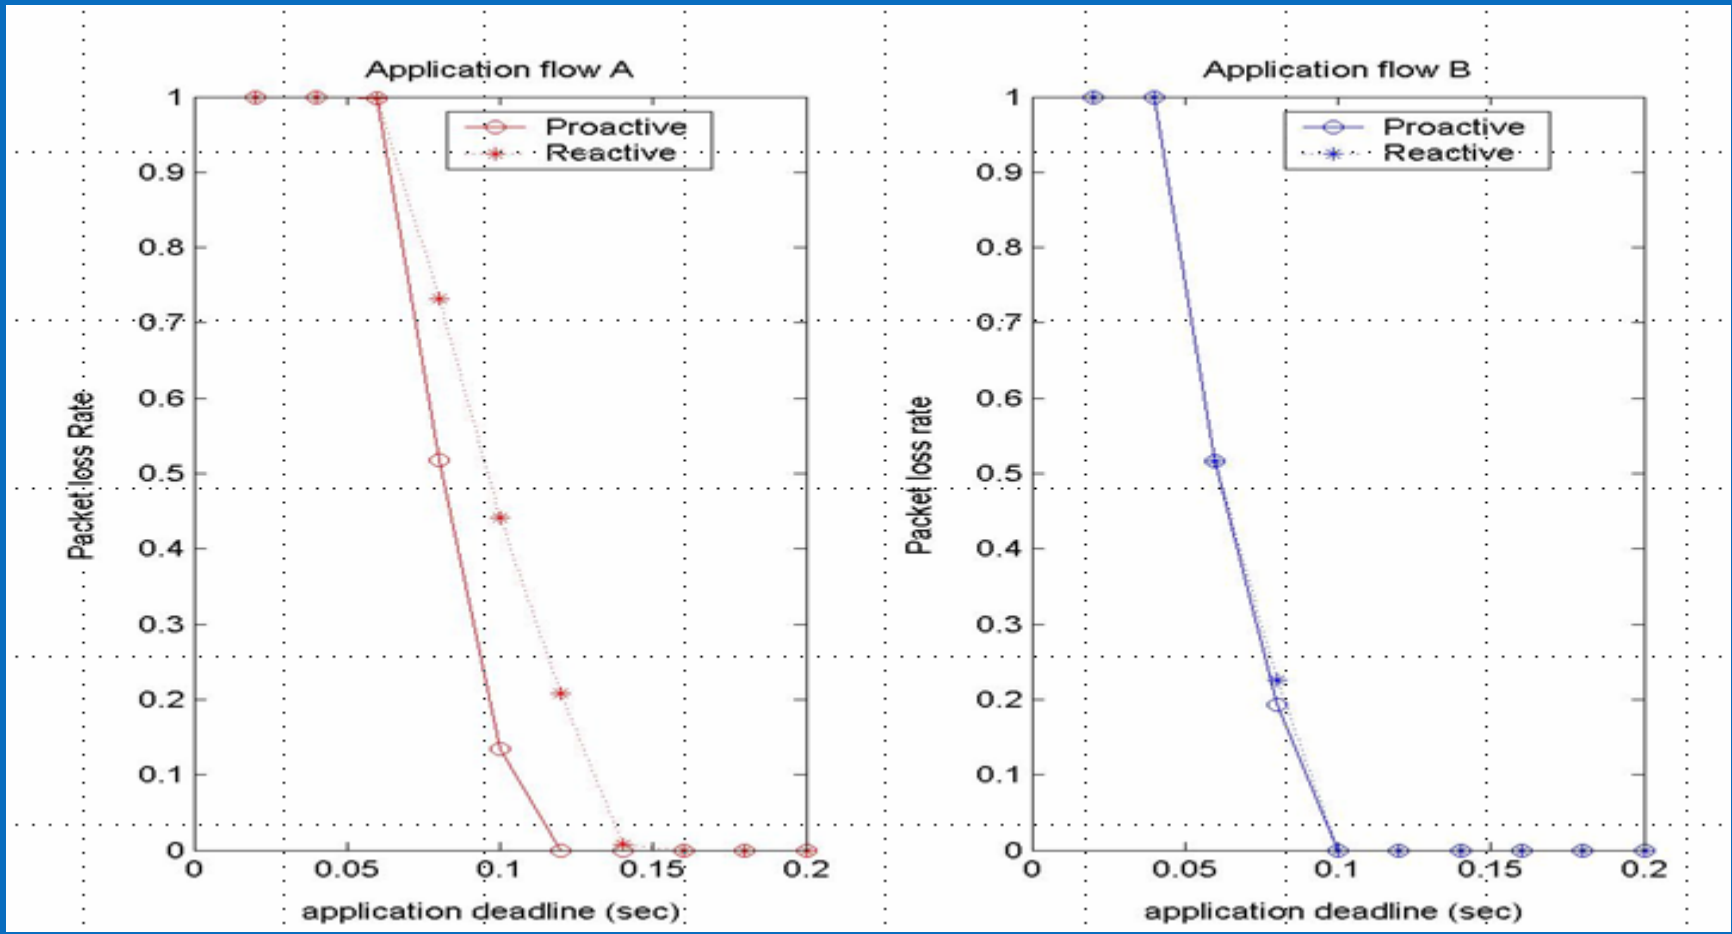

Generalized Cross-Layer Cross-Overlay Architecture

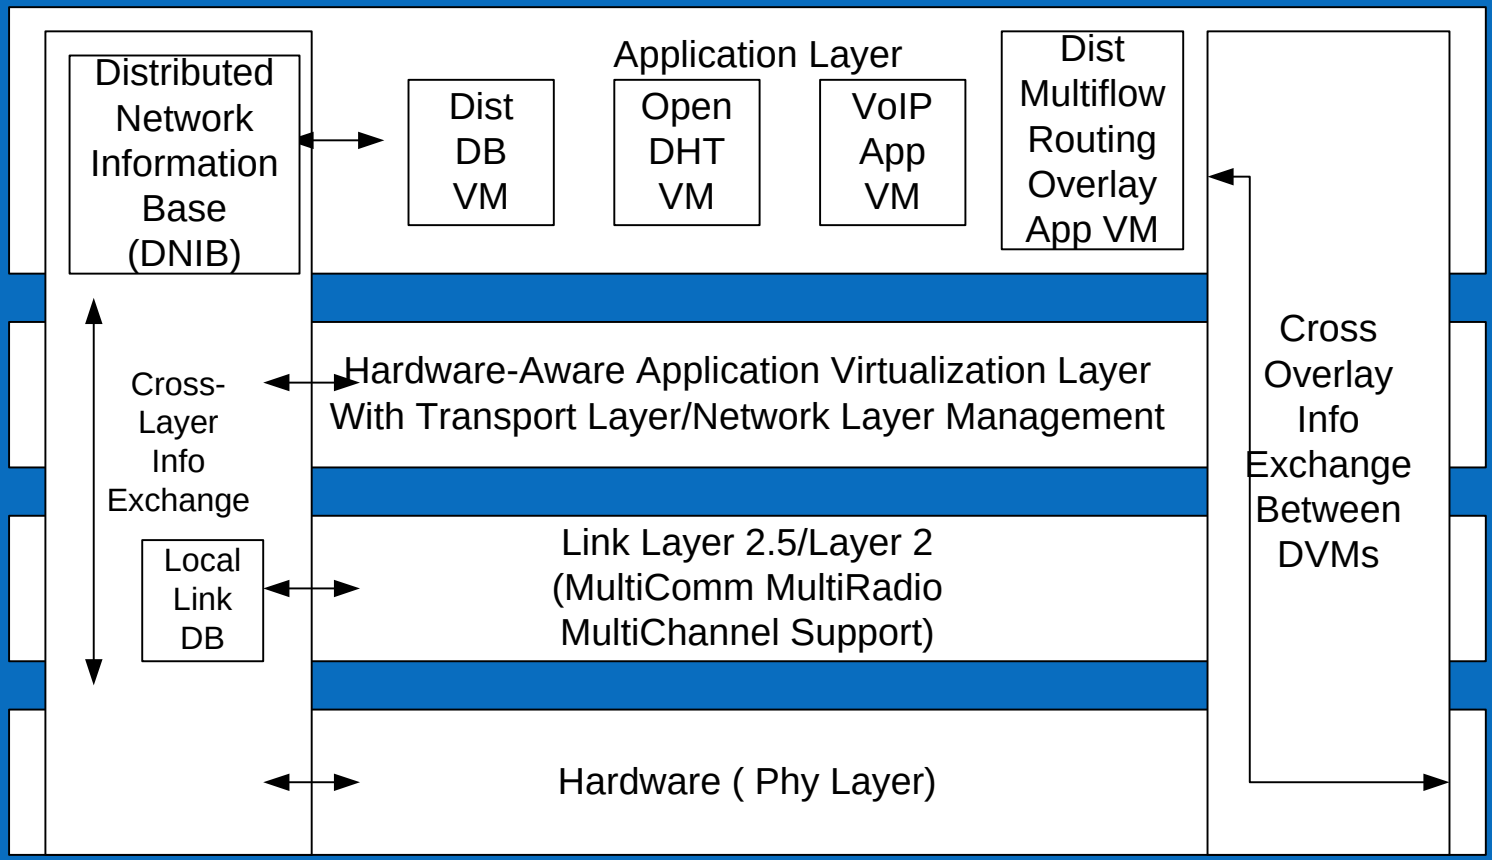

Example End-to-End Heterogenous Network Path (VoIP call)

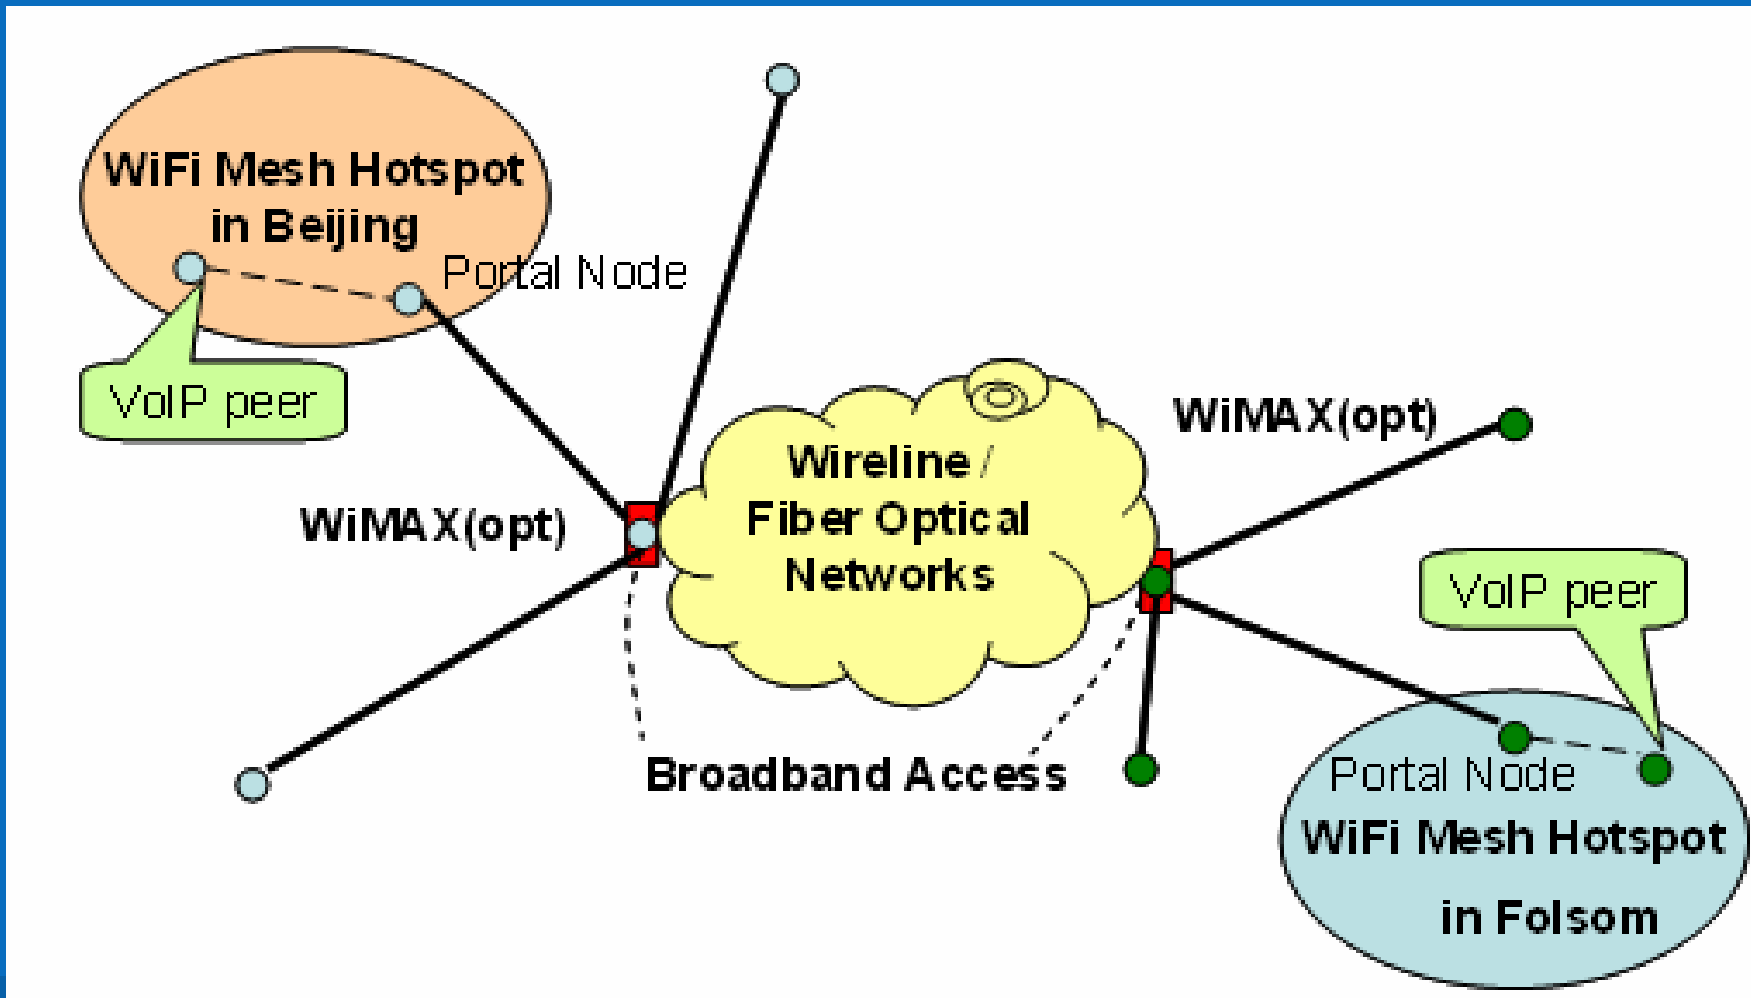

Questions?

intel


intel®
Leap ahead™