



High performance and efficient single-chip small cell base station SoC

Kin-Yip Liu

Cavium, Inc.

kliu@cavium.com

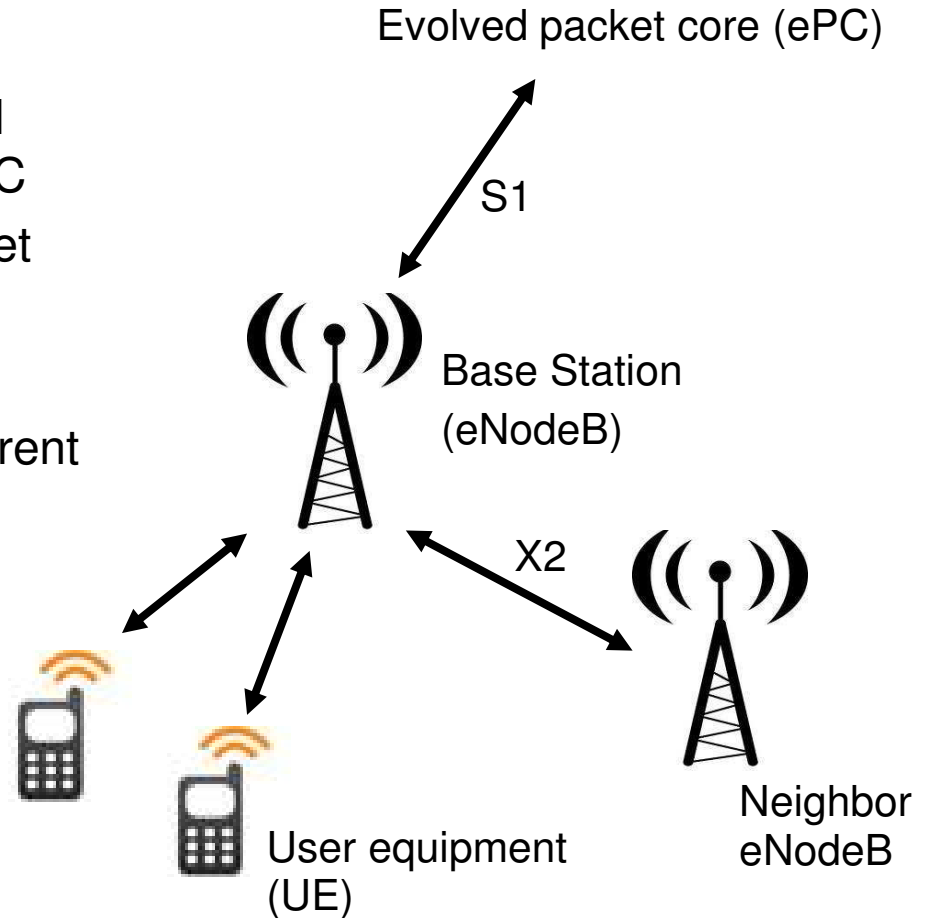
Hot Chips 24, August 2012

- Base station processing overview
- Why small cells and heterogeneous Radio Access Network (RAN)
- Small cell design based on OCTEON Fusion
- OCTEON Fusion CNF71XX architecture
- CNF71XX design
- Software models
- Summary

LTE Wireless Network Overview



- LTE equipment:
 - Base Stations – eNodeB
 - User equipment (UE), e.g. cell phone, dongle for notebook PC
 - Core network – Evolved Packet Core (ePC)
- An eNodeB interfaces with:
 - ePC (multiple nodes with different functions)
 - Control, signaling
 - To voice & data networks
 - UE's
 - Neighbor eNodeB's
 - Communicate load and interference info
 - Handover UE's



- eNodeB relays information between UE and ePC
- eNodeB and UE communication protocol:

Protocol layers	Processing functions
RRC (layer 3)	Set up and maintain radio bearers. Manage radio resources. Control functions. Handover decisions
PDCP (layer 2)	En/decrypt over-the-air traffic, Header de/compression
RLC (layer 2)	Segment and reassemble traffic. Ensure in-order traffic delivery. Re-transmit as needed
MAC (layer 2)	Schedule use of over-the-air resources. Select PHY configuration for transfers. Collect stats & report to RRC
PHY (layer 1)	Physical layer: OFDM for downlink. SC-FDMA for uplink

- eNodeB and ePC communication protocol:
 - IP network, IPsec protected, GTP tunnels of user data in UDP/IP, SCTP for control traffic

Classes of Base Stations

← Small Cells →

	Home Femto	Enterprise Femto	Pico	Micro	Macro
Cell Radius	50m	75m	250 - 400m	2 - 20km	20km
No. of users	8	32	128	1200	3600
Peak data rate	50Mbps DL 25Mbps UL	100Mbps DL 50Mbps UL	150Mbps DL 75Mbps UL	300Mbps DL 150Mbps UL	900Mbps DL 450Mbps UL
User Mobility	4 km/hr	4 km/hr	50 km/hr	350 km/hr	350 km/hr
Locations	Home	Office, school, apartment buildings, malls	Urban hotspots, rural areas	Urban, rural areas	Metro, traditional approach


DL – Downlink. Traffic going from network to user
 UL – Uplink. Traffic going from user to network

Additional Small Cell Requirements


- WiFi option
 - Single platform for Small Cell + Access Point
 - SoC must provide performance headroom for both functions
- Power-over-Ethernet
 - Simplify system deployment, but limited system power supply
 - SoC must consume very low power
- Time synchronization
 - Mandatory for LTE base stations. IP backhaul, no TDM interface
 - GPS option. May not work well in-door
 - Software solutions: IEEE 1588 v2, NTP. In-door OK, cost effective
- Security
 - Authenticated and encrypted software for secure boot

Why deploy small cells?

.....for **Hot spots** and **Not spots**




Easing congestion
within macro coverage



New coverage
in addition to macro

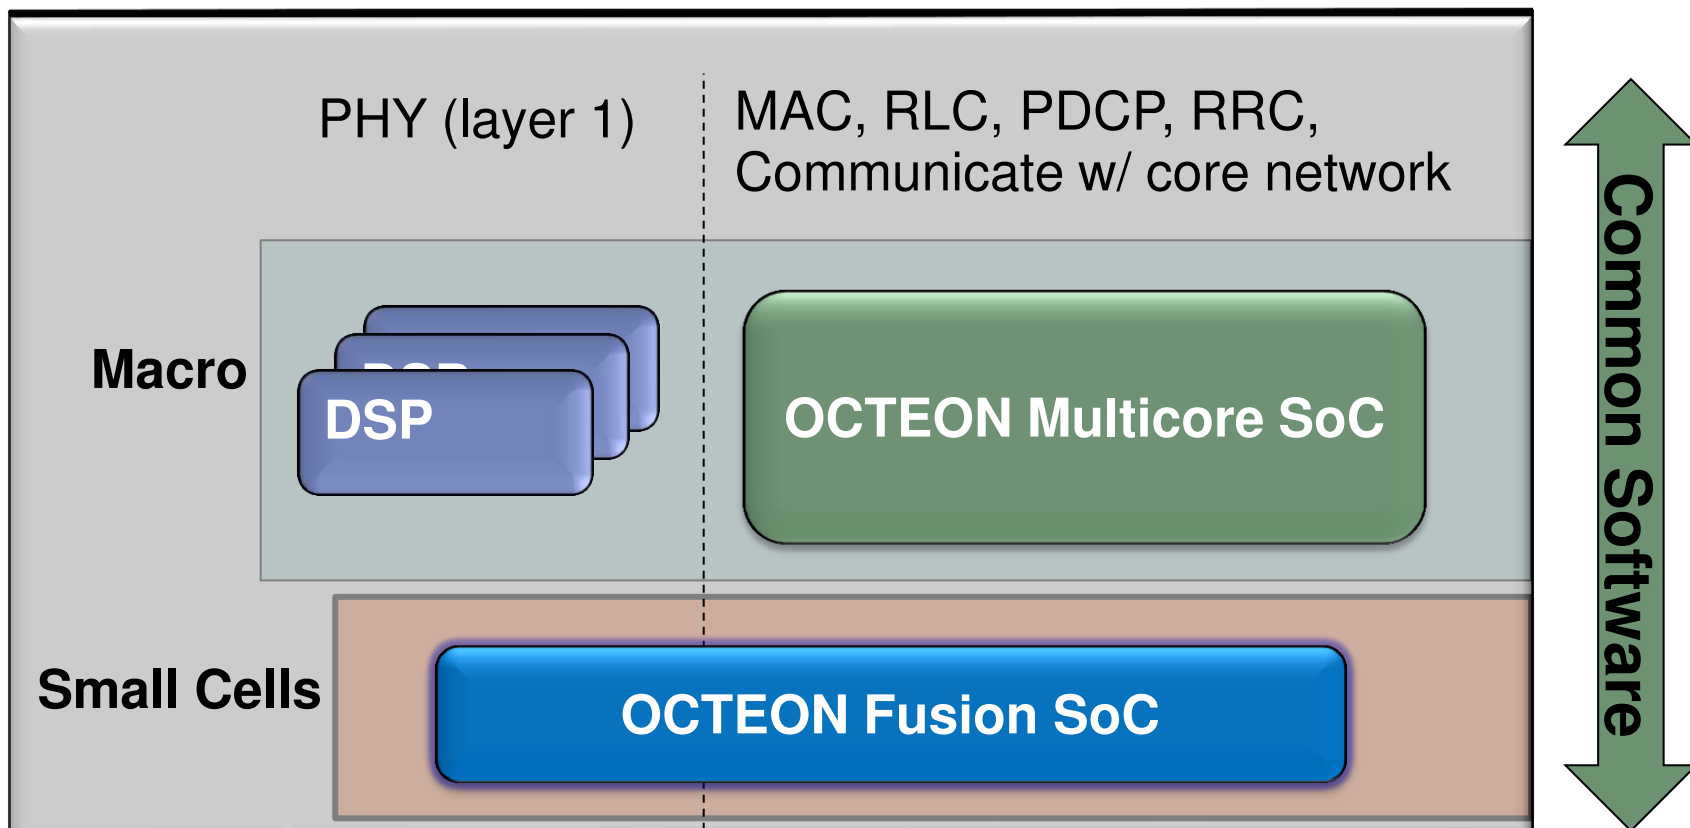
**Small Cells essential for LTE
coverage, capacity, and throughput**

Current Generation Base Stations




Single-chip Multicore SoC for Layer 2 and above processing. Common software from Small to Macro cells

Next Generation Base Stations



Single-chip Multicore + baseband module SoC for Small Cells. Common software from Small to Macro cells



Small Cell Base Station + Access Point

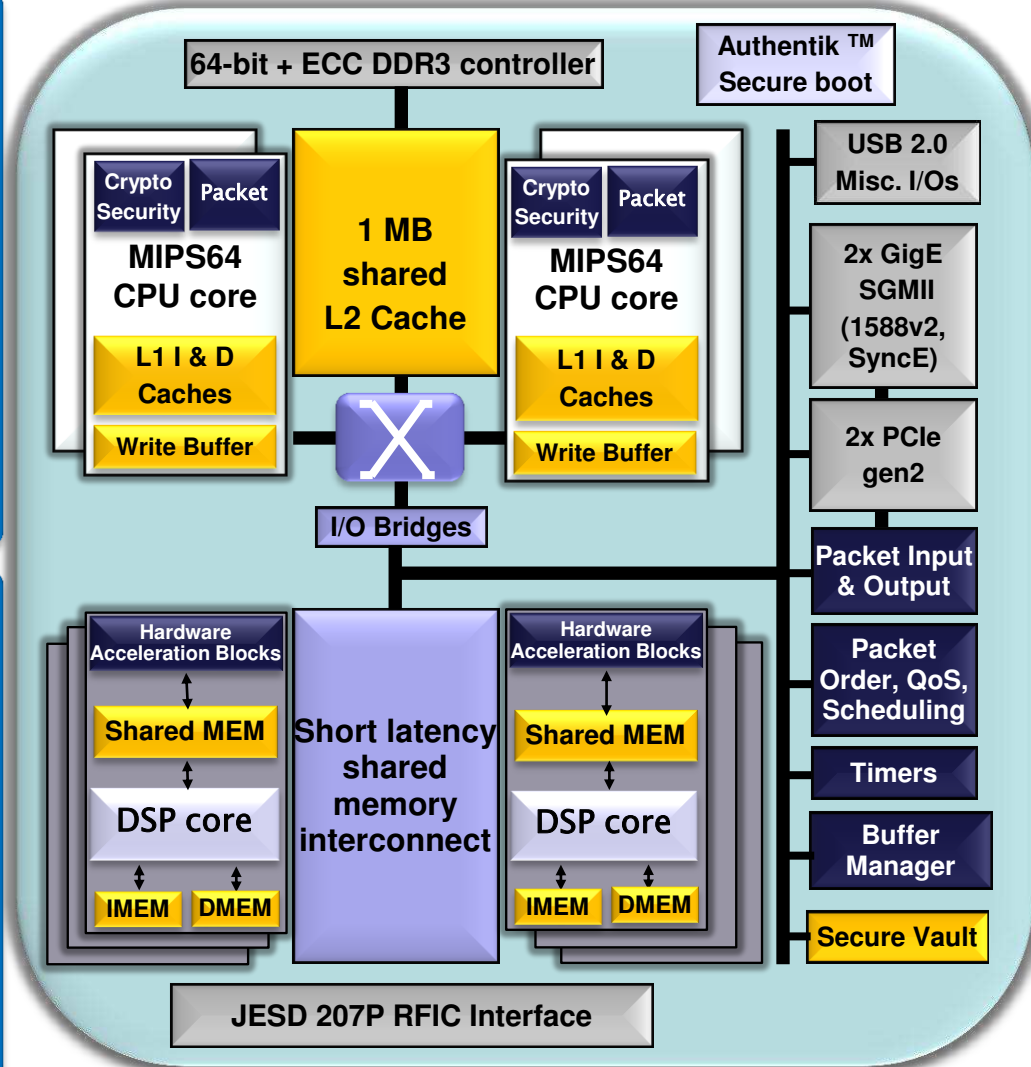
OCTEON Fusion CNF71XX

Small cell BaseStation-on-a-chip Family



OCTEON MULTICORE

BASEBAND MODULE



LTE TDD/FDD, WCDMA
2x2 MIMO, Up to 20 MHz

- **High Performance LTE / 3G Small Cell SoC Processors:**
 - 4 MIPS64 cores up to 1.5 GHz
 - 6 DSP cores up to 500MHz
 - Many HW Accelerators for Packet Processing, LTE/3G, and Security
 - IEEE 1588 v2, SyncE
 - Authentik secure boot
- **Highly Scalable**
 - Spanning 32 to 200+ Users
 - 3G and LTE FDD & TDD
 - Up to LTE 20MHz 150 Mbps Uplink (UL) + 150Mbps Downlink (DL)
- **Headroom for Unique Carrier Class Features**
 - Multi-User MIMO
 - Self Optimizing Networks
 - Interference Cancellation
 - Advanced Receivers

Design Philosophy



High Performance and Power Efficient

- Power and area efficient CPU and DSP cores
- Scale performance with more cores
- Not depend on very high frequency or core complexity

Short Latencies Deterministic Performance

- Shortest cache and memory latencies. Optimize for determinism
- Flexible prefetch, cache hints, options to cache packet headers only
- L2 way partition feature avoids cache pollution

Optimized ISA Ease of programming

- MIPS64 r3 instruction set + >80 OCTEON instructions
- Full C programming. Standard OS and development tools

Comprehensive Hardware Acceleration

- TCP/IP, complete packet receive and transmit offload, packet ordering, QoS, work scheduling, buffer de/allocation, IPSec, wireless crypto algorithms, timers, wireless baseband functions
- Crypto coprocessor in each core. Best latency & determinism

Software Compatible Roadmap

- Software compatible from 1-48 cores and across generations
- Single SDK to develop software for all OCTEONS
- Software for macro base stations directly reusable for Small Cells

Baseband module processing flows

- Wireless UL and DL processing differ. Partition the DSP cores and assign relevant hardware accelerators for UL Vs. DL processing
- Modular design with flexible partitioning simplifies software design

6x DSP cores optimized for wireless baseband processing

- 3-way VLIW, with 16x MAC or 4x complex MAC vector processing per cycle
- Optimizing instructions for wireless baseband processing
- Dual 128-bit load/store paths transfer up to two vector operands each cycle


Hardware accelerators (HABs)

- Comprehensive set of LTE and 3G, UL and DL relevant accelerators
- Automate offload to accelerators with DMA engines and Sequencer

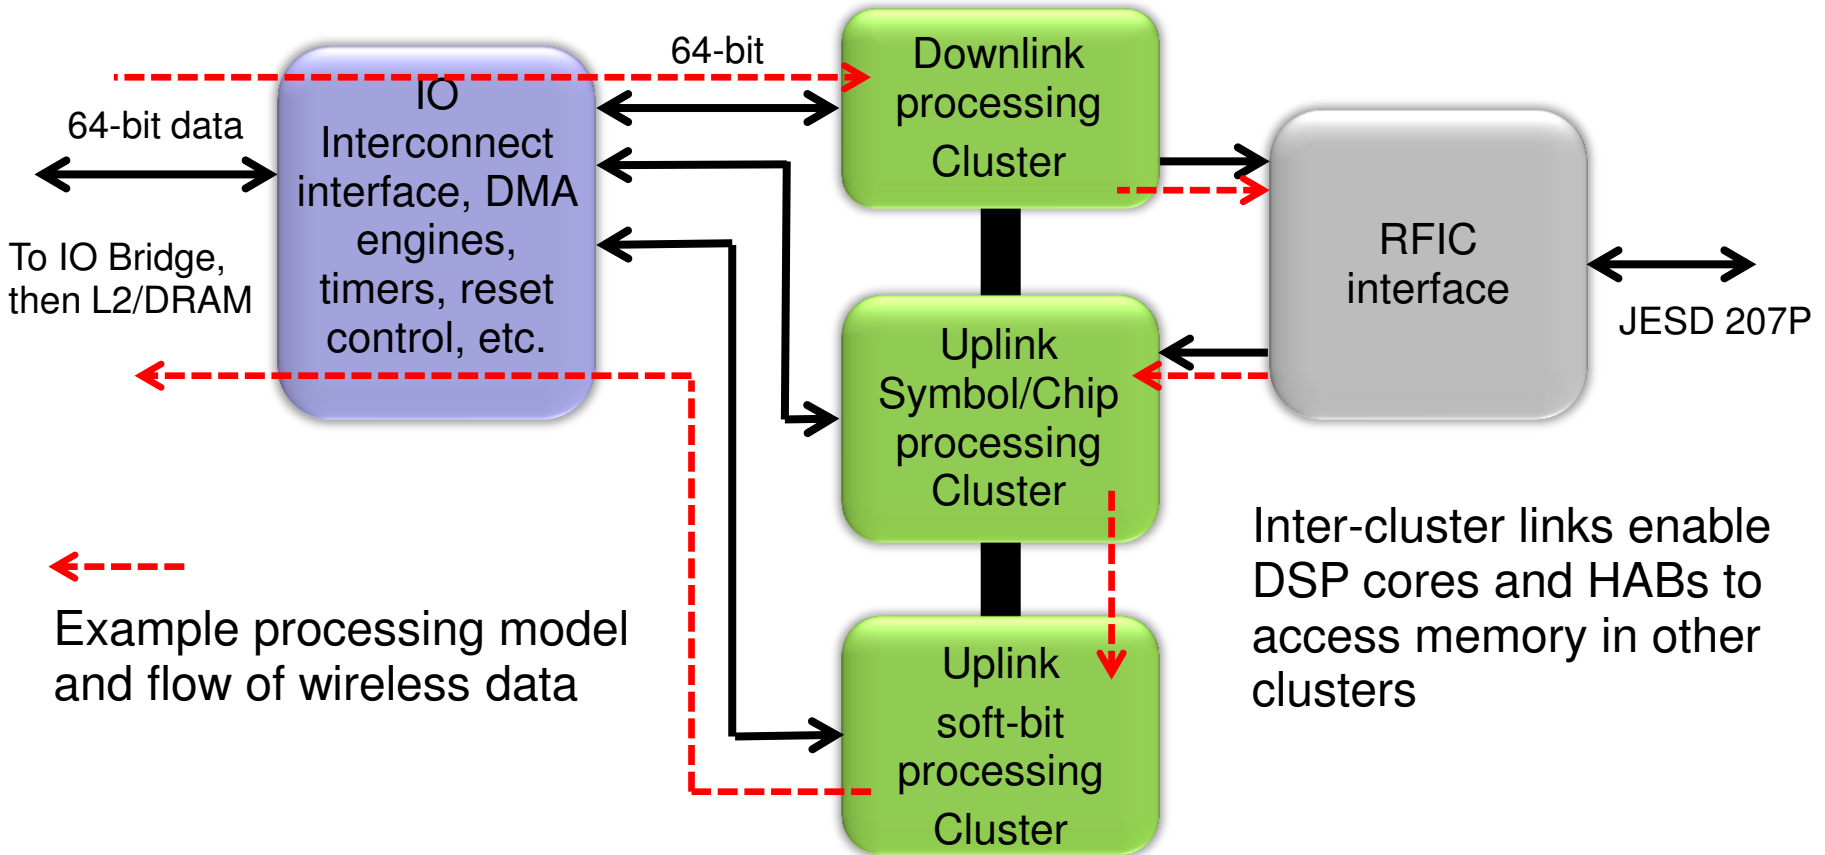
Shared memory interconnect

- DSPs and HABs can access any memory structure in entire baseband module

A Cluster of the Baseband Module



128-bit dual load/store paths enable VLIW DSP cores to fetch two 128-bit vector operands + processing in single cycle



Shared memory interconnect enables flexibility in optimizing the processing models and flows

Wireless L2 & L3, Transport, Control, WiFi, Customer Apps

- OCTEON Fusion = OCTEON Multicore + Baseband module
- The OCTEON Multicore part of the SoC is the same architecture as OCTEON Multicore SoCs which have been widely deployed for designing base stations

CPU cores

- 4x OCTEON MIPS64 cores
- Shortest L1 and last-level-cache (L2) latencies among multicore processors
- Power optimizer™ per-core software controlled power reduction
- Fine-grained clock gating

Hardware accelerators

- Comprehensive packet processing hardware: Headers parsing, classification, RED, QoS, buffer allocation, L4 checksums, traffic rate limiting & scheduling
- Crypto, packet order, work scheduling, timers for TCP and RLC, RoHC

Low latency interconnect


- Split-transaction interconnects and L2 cache run at core frequency

Custom designed efficient 64-bit CPU core

- Dual-issue, 8+ stages. Optimized for perf/watt, perf/area
- Short 3 cycles L1 cache load-to-use latency
- MIPS64 r3 instruction set + >80 optimizing instructions

Examples of optimizing instructions added

- Atomic memory ops (increment, add, fetch-and-add, etc.)
- Insert/extract arbitrary bit fields within a word
- Branch if certain bit field contains a set bit or not
- Compare operands and set bit0 for equal / not equal
- Additional flavors of prefetch and cache hints
- Population count
- Unaligned load/store



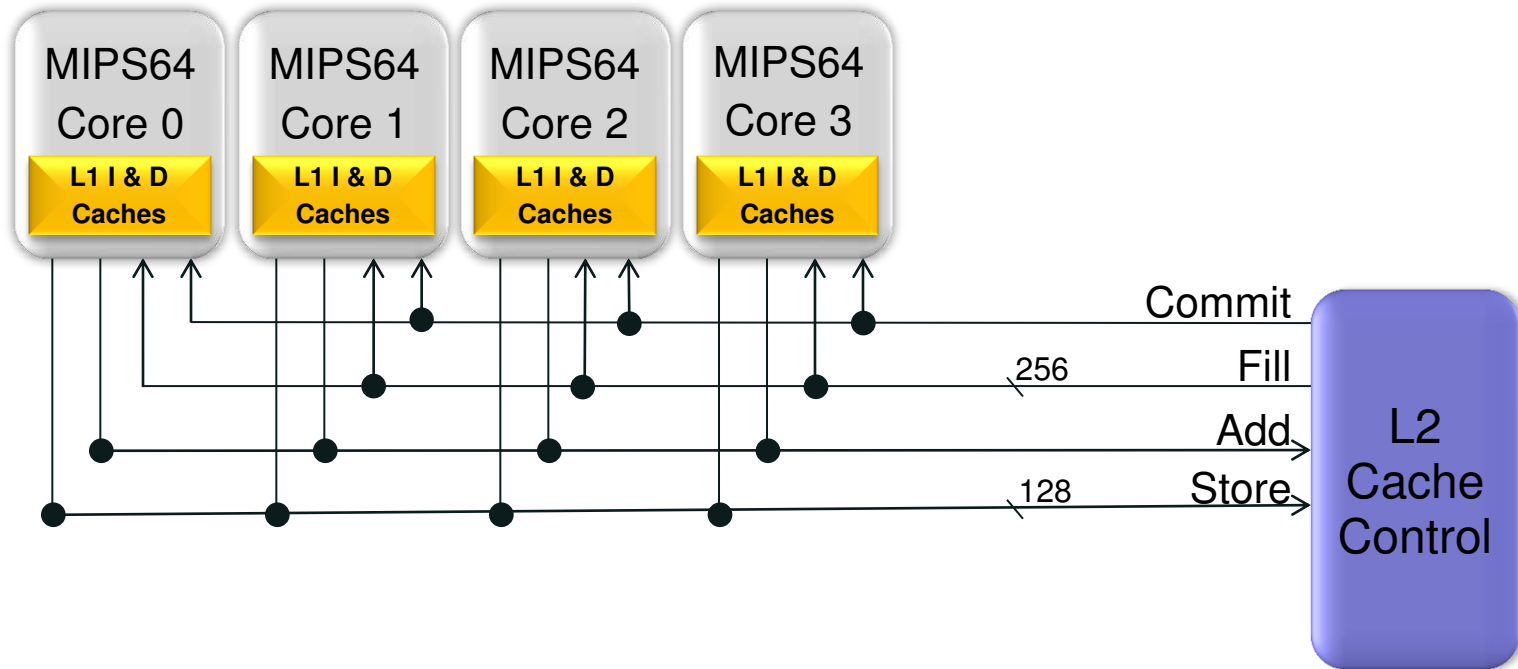
L1 <-> L2 Cache: Write-through

- Excellent performance for networking and wireless applications
- Minimal per-CPU-core cost (power, area)
- Lowest possible read latencies
- Allows many outstanding stores, optimizations
- Automatic L1 error correction

L2 Cache <-> DRAM: Write-back


- Standard DDR3 DRAM DIMM's are highest performance with block transfers
- Minimizes required DRAM bandwidth
- Don't-write-back feature (e.g. for most of packet data) plus additional cache hints

CNF71XX Coherent Interconnect

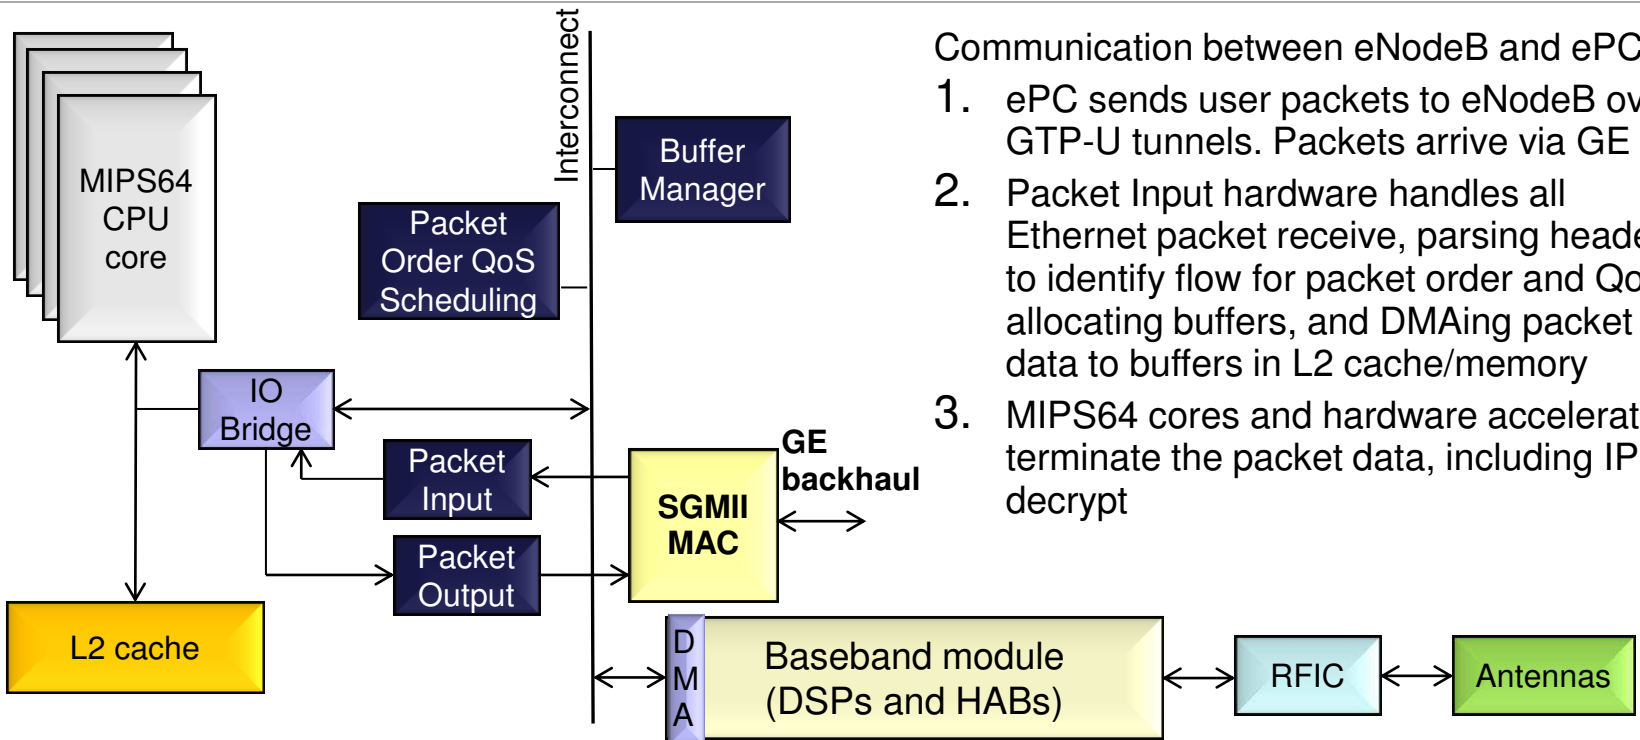


64-bit CPU cores, split-transaction interconnect,
L2 cache & controller all run at core frequency

CNF71XX Chip Floorplan



Packet/Data Flow: LTE Downlink (DL) Processing



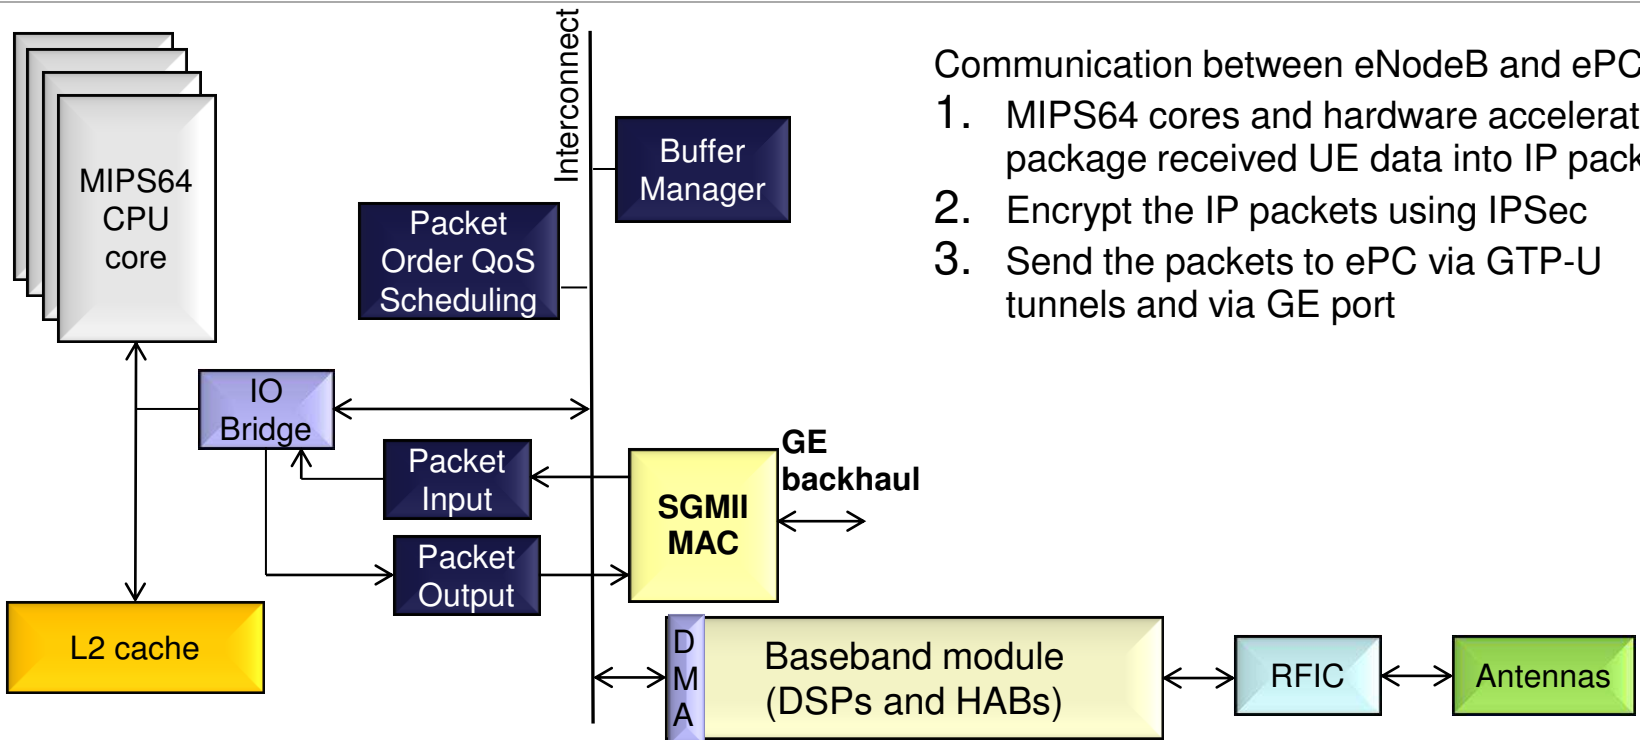
Communication between eNodeB and ePC:

1. ePC sends user packets to eNodeB over GTP-U tunnels. Packets arrive via GE port
2. Packet Input hardware handles all Ethernet packet receive, parsing headers to identify flow for packet order and QoS, allocating buffers, and DMAing packet data to buffers in L2 cache/memory
3. MIPS64 cores and hardware accelerators terminate the packet data, including IPsec decrypt

Communication between eNodeB and UE's with 1ms TTI (transmission time interval):

1. MIPS64 cores and accelerators process PDCP, RLC and MAC protocol layers.
2. MAC layer processing schedules data and wireless PHY configuration for DL transmission
3. Baseband hardware DMA's data from L2 cache to its local memory
4. Downlink DSP cores and HABs complete DL processing and transmit data out via RF interface

Packet/Data Flow: LTE Uplink (UL) Processing



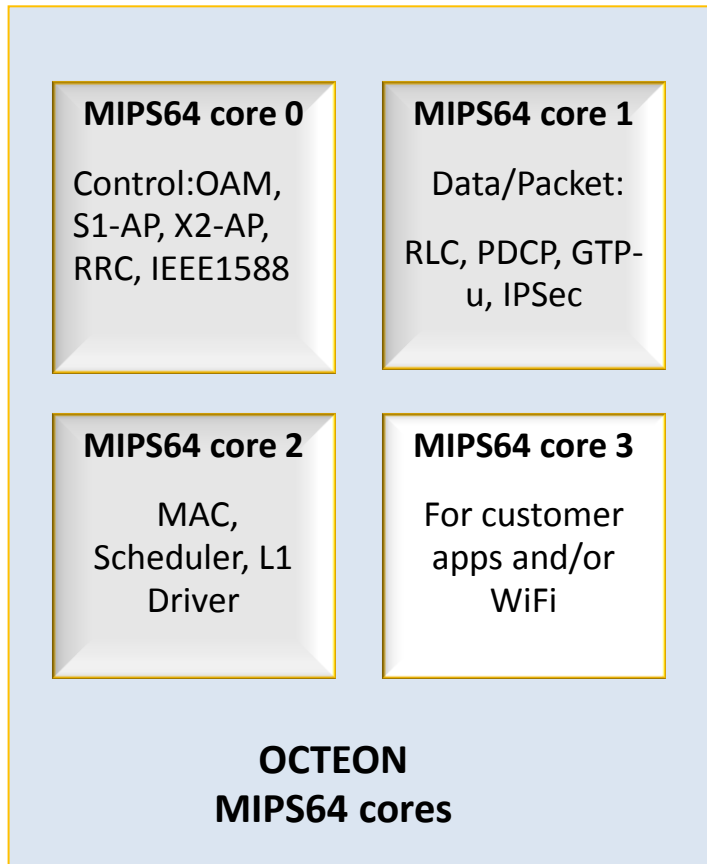
Communication between eNodeB and ePC:

1. MIPS64 cores and hardware accelerators package received UE data into IP packets
2. Encrypt the IP packets using IPsec
3. Send the packets to ePC via GTP-U tunnels and via GE port

Communication between eNodeB and UE's with 1ms TTI (transmission time interval):

1. PHY baseband processes UL traffic and detects random access from UE's
2. PHY baseband DMAs processed UL data to L2 cache
3. MIPS64 cores and accelerators process MAC, RLC, and PDCP layers to terminate received UE traffic into packets.

Mapping eNodeB to Multicore



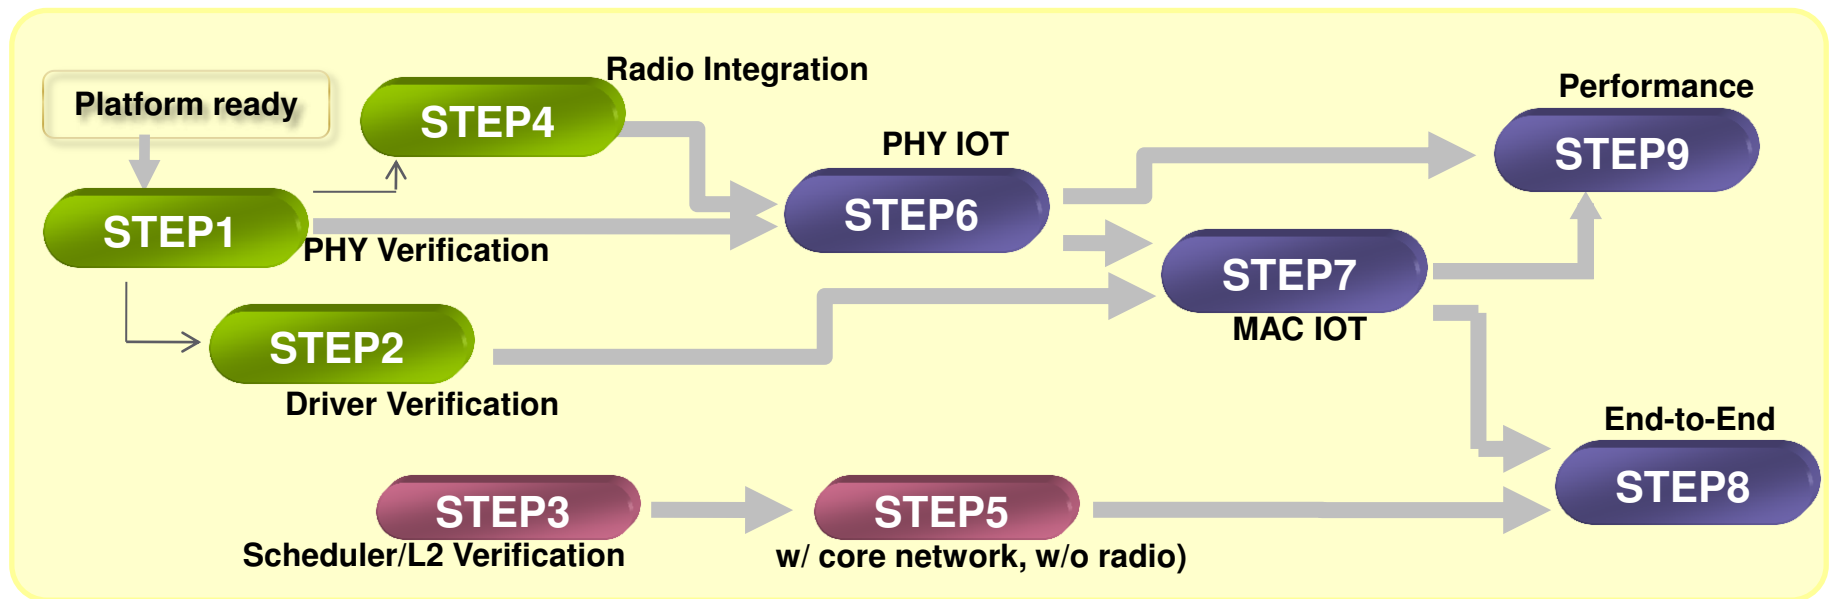
- Example partitioning : LTE eNodeB AP
 - MAC and L1 driver on one core
 - Easy to meet LTE 1ms TTI
 - Quick response to PHY interrupts
 - RLC, PDCP, Transport on one core
 - Option to partition L2 cache to avoid cache pollution from control processing
 - Control processing on one core
 - 1 core free
 - Headroom for WiFi and service provider applications
- Small Cell Forum API compliant

Quad-core delivers required headroom and deterministic performance for real-time LTE and other processing

CNF71xx Complete End-to-end Validation



- › **STEP1 – PHY + Driver S/W + PLT (Physical Layer Test)**
- › **STEP2 – PHY + Driver S/W + Scheduler**
- › **STEP3 – L1 + L2 + L3**
- › **STEP4 – PHY + Modem + Radio**
- › **STEP5 – Core network + Basestation (L2/L3 stacks, S1 I/F)**
- › **STEP6 – IOT (Interoperability Testing) in PHY (PLT + Modem + Radio + UE L1)**
- › **STEP7 – IOT in MAC (w/ UE L1/L2)**
- › **STEP8 – IOT in E2E (w/ UE over full protocol stacks)**
- › **STEP9 – DL/UL Performance Measurements w/ UE**



➤ OCTEON Fusion CNF71XX

- High performance “base station on a chip” SoC
 - LTE 20MHz, 150Mbps DL + 150Mbps UL, 2x2 MIMO, 128 users
- OCTEON Fusion = OCTEON multicore + baseband
 - Same OCTEON software for small to macro cells
- End-to-end interoperability and performance verified

➤ Optimized for Base station designs

- Delivers deterministic real-time performance, low power, and high integration, with significant compute headroom
 - 4x enhanced & efficient 64-bit (OCTEON MIPS) CPU cores
 - 6x Baseband optimized DSP vector processors
 - Many hardware accelerators
 - Optimized for short latencies and deterministic performance

Backup



Cavium: Company Summary















- **Founded** 2001
 - **NASDAQ IPO (CAVM)** 2007
 - **Locations:** US, India, China, TW
 - **2011 Revenues :** \$259M, +26% YOY
 - **5 year CAGR:** ~50%
 - **Profitable with Strong Financials, Zero Debt**
-
- **Addressing Multi-billion dollar Networking, Communications, Storage and Digital Home markets.**
 - **MIPS64 and ARM based Multi-core Processor SoCs; Multi-core Search and Security Processors**
 - **All Top Networking, Wireless and Security Vendors use Cavium**

Carriers coping with 1000x traffic increase and no extra revenue



Smart devices multiply traffic

Smartphone		=	 x 24*
Handheld Gaming Console		=	 x 60*
Tablet		=	 x 122*
Mobile Phone Projector		=	 x 300*
Laptop		=	 x 515*



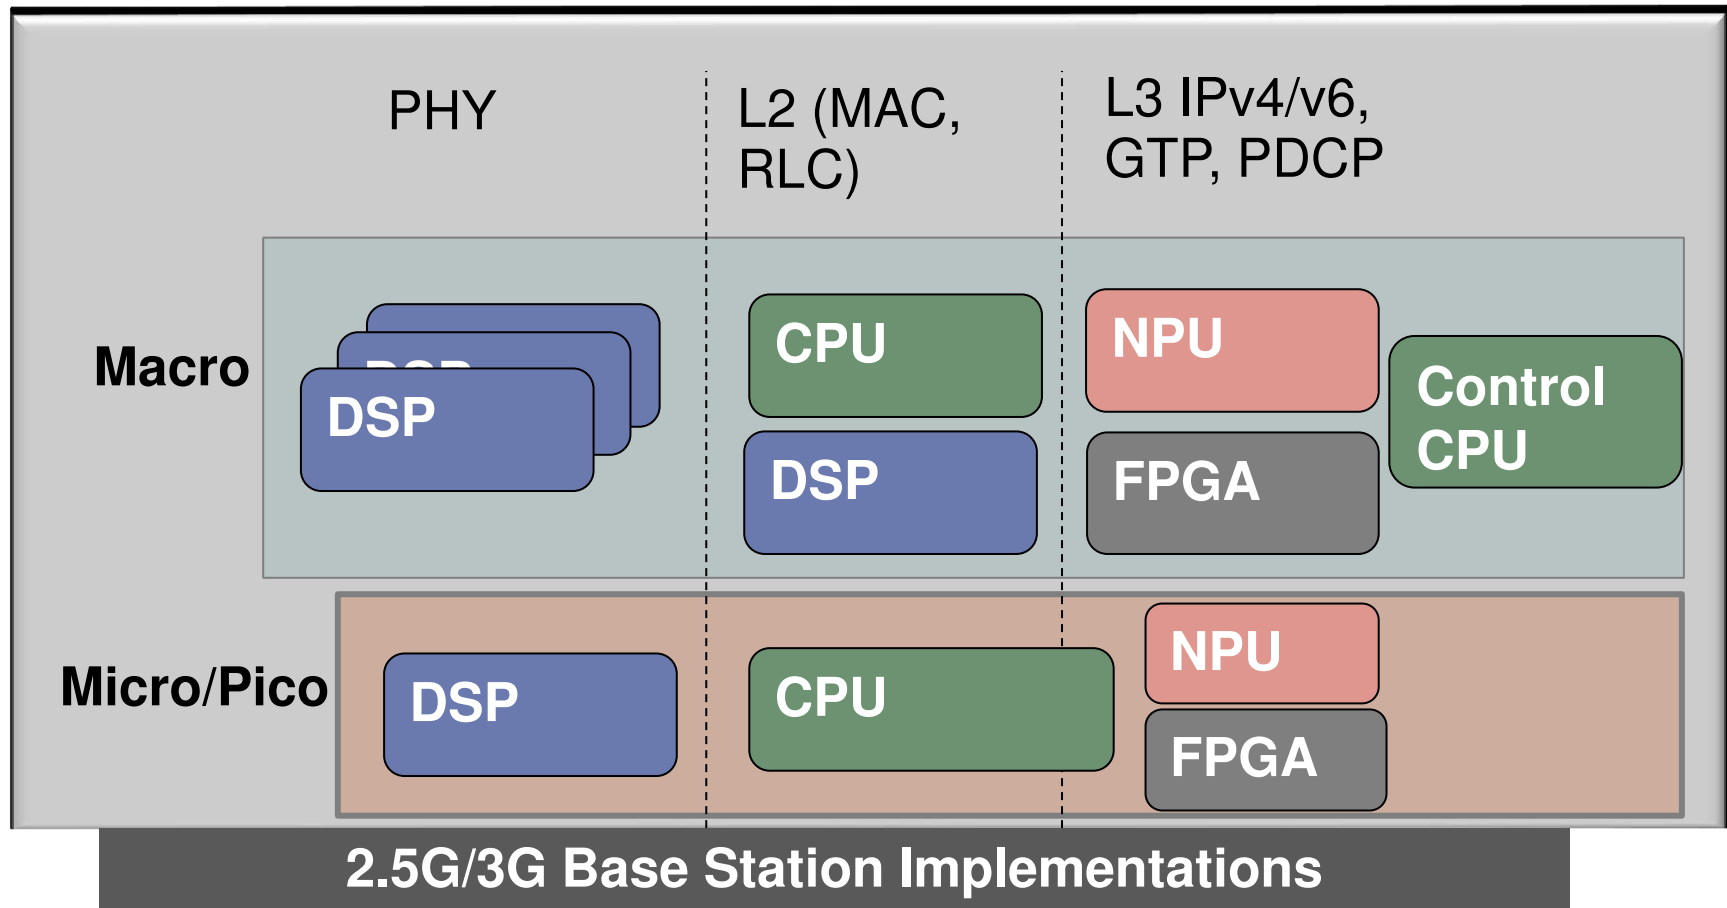
* Monthly basic mobile phone data traffic

Source: Cisco VNI Mobile, 2011

Heterogeneous Radio Access Network

- Macro base stations are expensive (CAPEX and OPEX)
- Augment Macro with Small cell base stations to add capacity and coverage cost effectively

Previous Generation Base Stations



Before Multi-core SoCs became available, Base Station designs required many components, microcode programming on NPU, general purpose CPUs, FPGAs, and many development environments. High complexity