

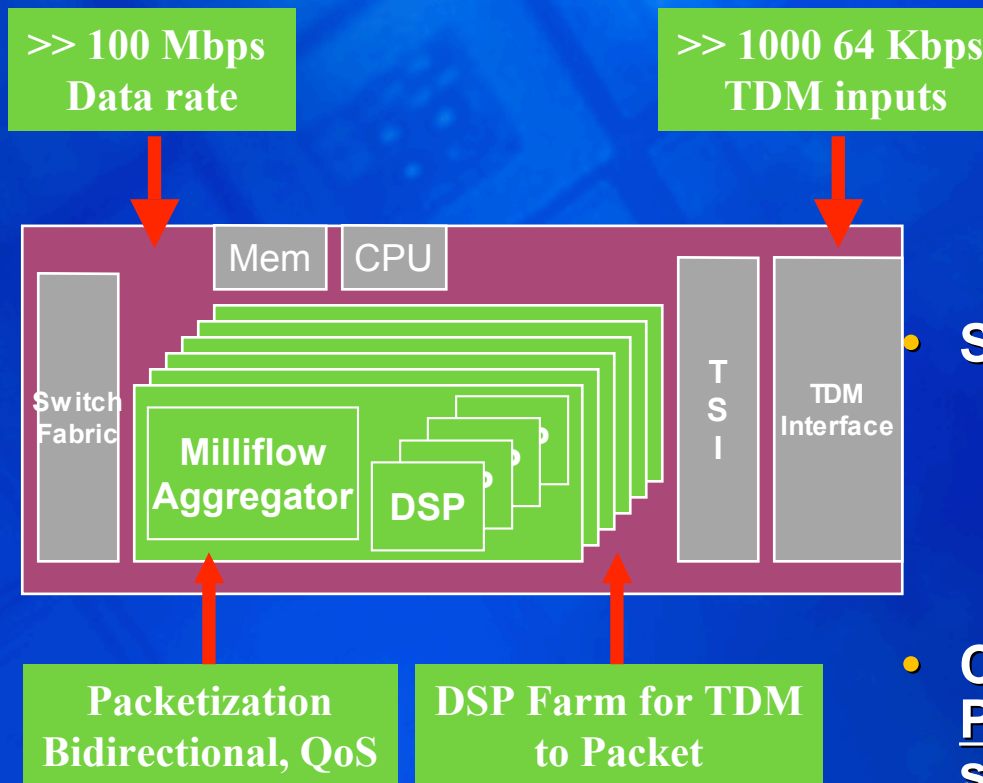
A Milliflow Aggregation Processor

Bapi Vinnakota

Objectives

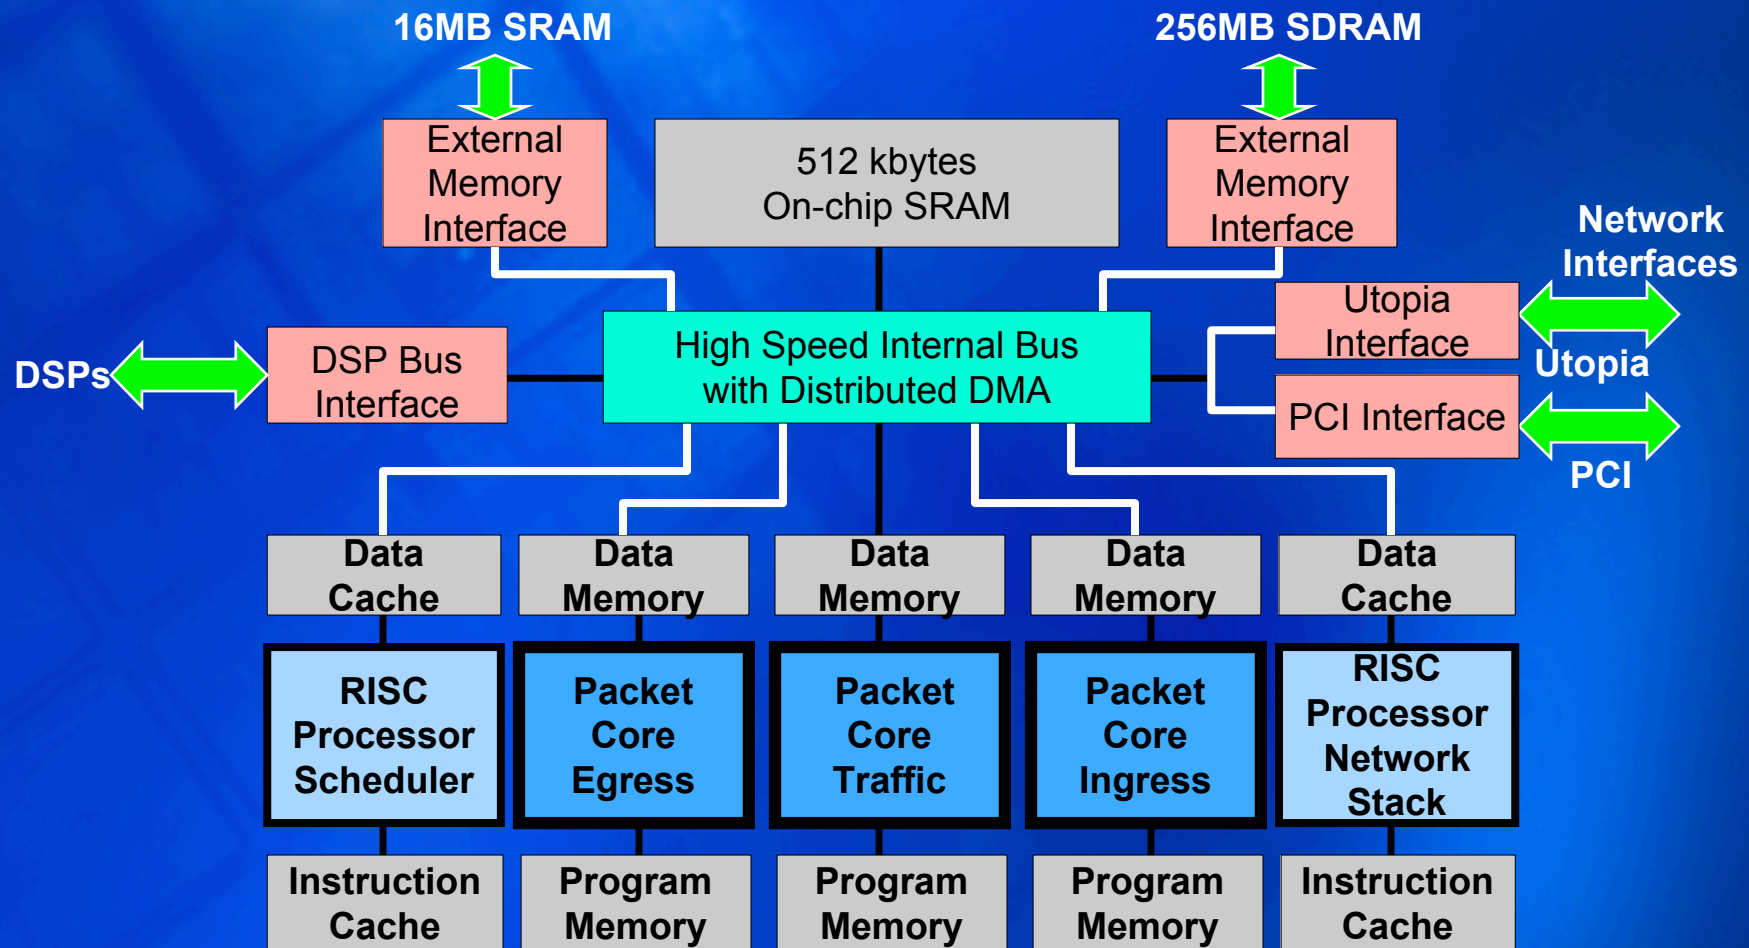
- **Milliflow aggregation systems**
 - Voice and other edge systems
- **SoC architecture**
 - High-level features
 - System dataflow
- **Component architecture**
 - Custom cores
 - Packet interfaces
- **Results**
 - Physicals
 - Performance
- **Conclusion**

Milliflow Aggregation Systems



- Milliflow aggregation nodes
 - Aggregate thousands of *milliflows*
 - The bandwidth of a milliflow is less than one thousandth of the bandwidth of the aggregate bandwidth
- Stringent constraints
 - Real time flows
 - Small packet sizes
 - System power dissipation is limited
- Custom Milliflow Aggregation Processor (Magpie) for VoIP systems
 - Experimental prototype chip, not designed for production

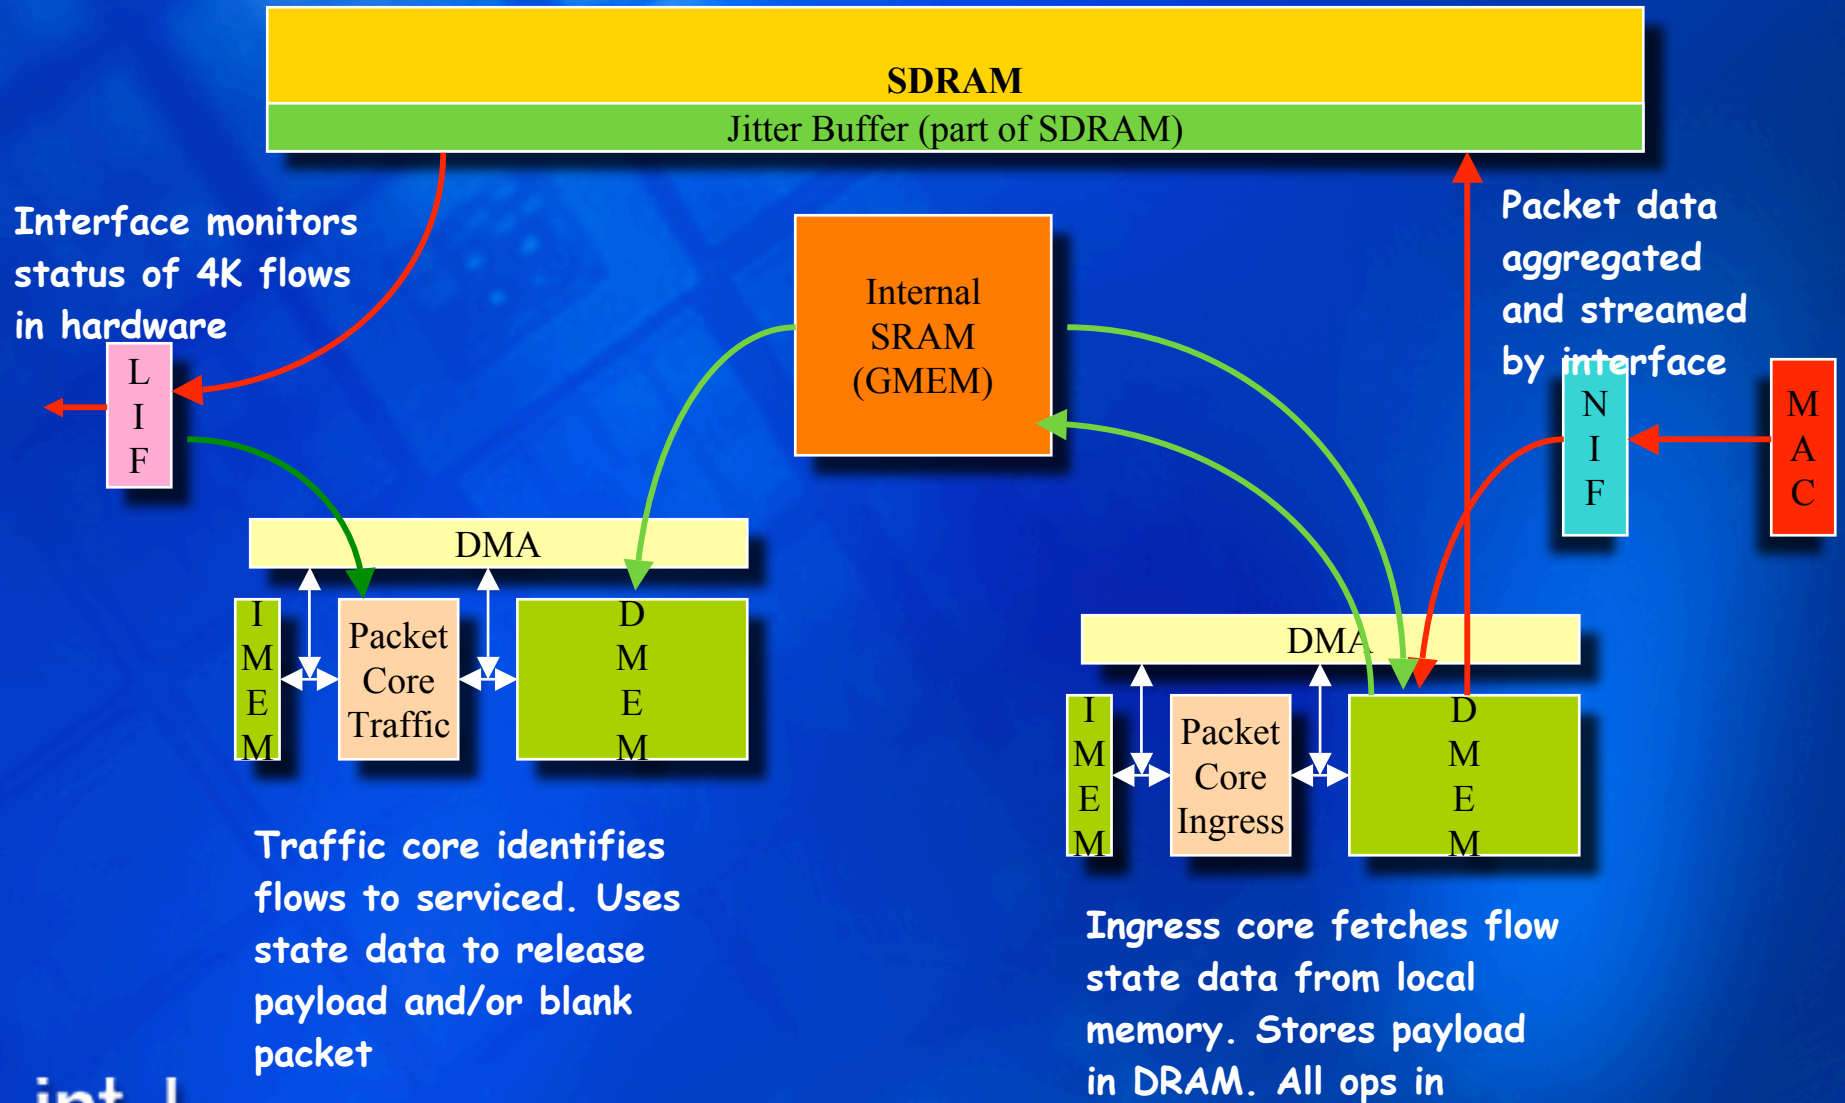
Magpie Architecture



Magpie High-Level Description


- **Key Functionality**
 - VoIP (RFC1489), VoATM (I366.2), VoIPoATM (RFC1483)
 - RTP(RTCP)/UDP/IP, RTP/AAL2 muxing, AAL5, TCP/IP
 - Jitter buffer algorithms and statistics support
- **Performance**
 - OC-3 5-ms, 2*OC-3 10 ms, OC-12 20 ms
- **Functional Interfaces**
 - 32-bit 66 MHz PCI (v.2.1)
 - 16-bit 50 MHz Utopia Level II
 - 32-bit 33 MHz VxBus (functionally glueless interface to IXS farm)
- **Memory Interfaces**
 - 64-bit SDRAM interface up to 256MB
 - 32-bit ZBT SSRAM interface up to 16MB
- **Architecture**
 - Three packet engines, two control cores, DMA engines

Packet Interface to DSP Farm Data Packet Flow





Interface monitors status of 4K flows in hardware

L I F



Traffic core identifies flows to serviced. Uses state data to release payload and/or blank packet



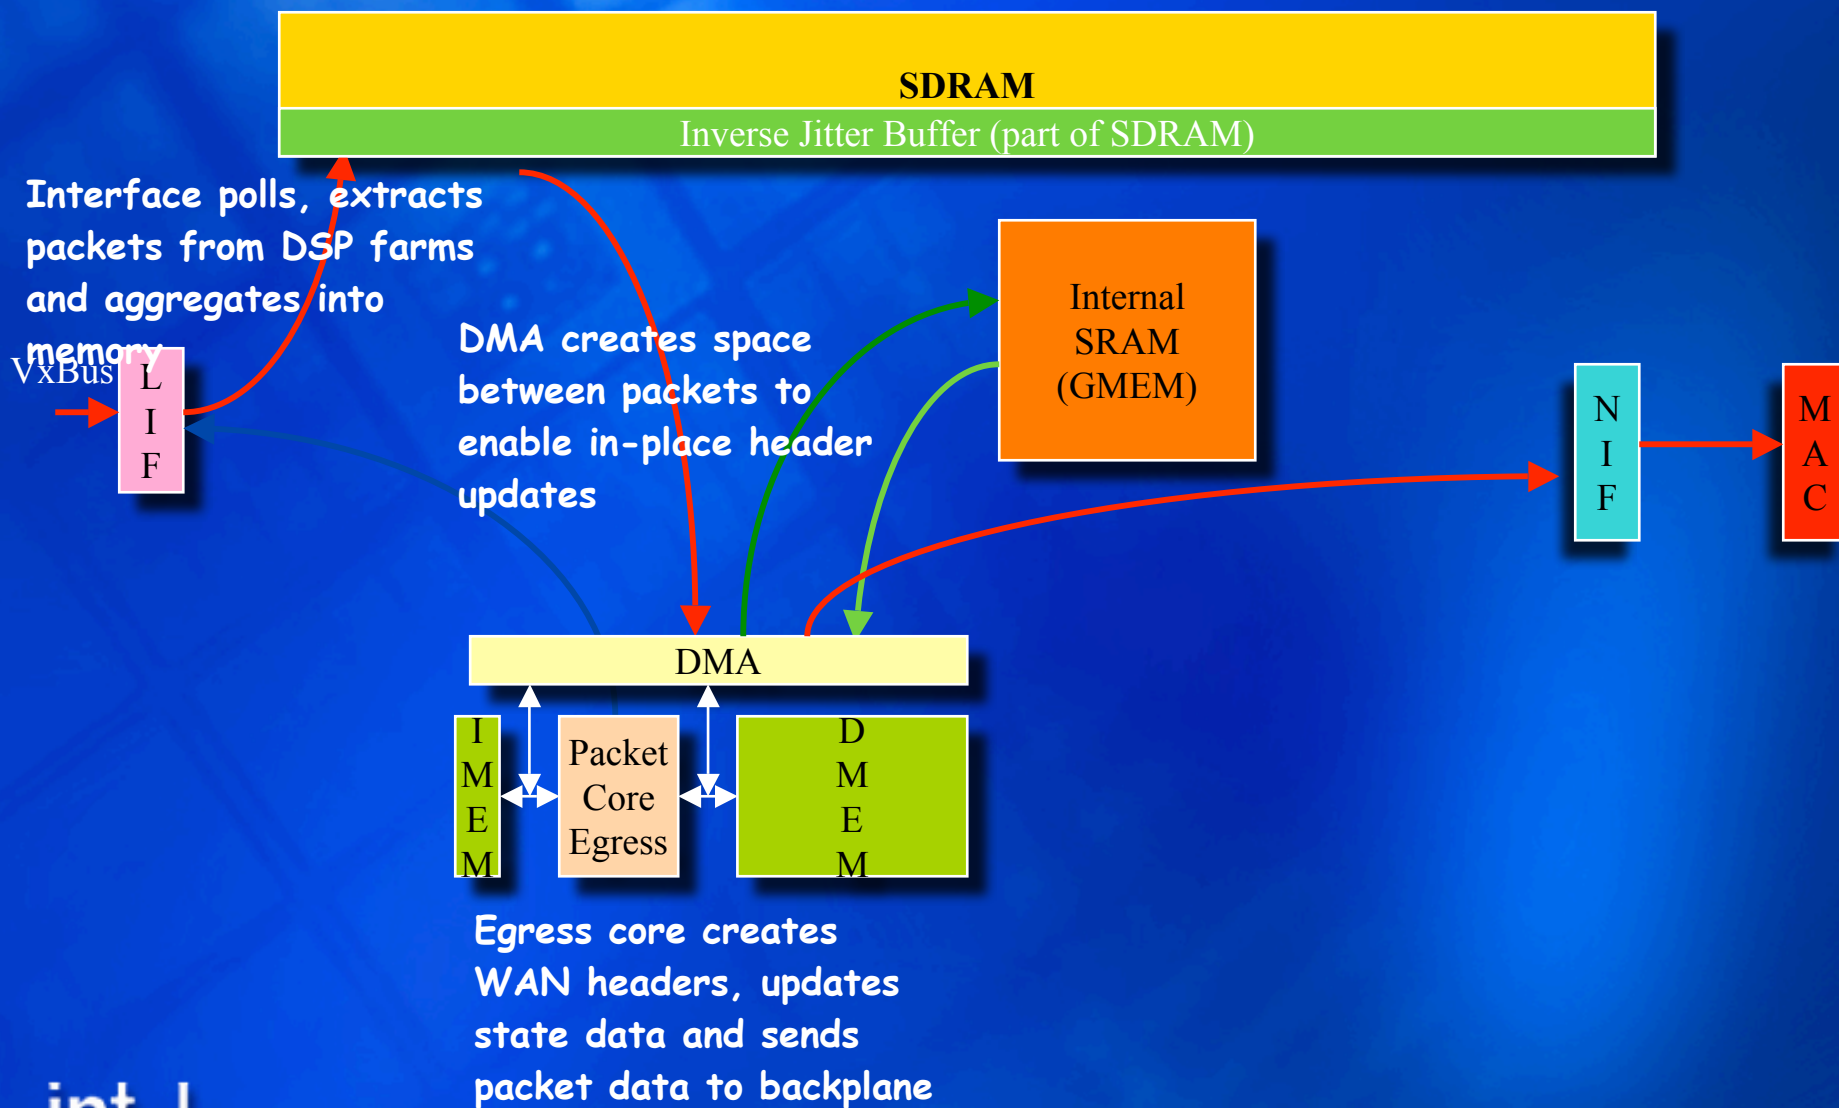
Ingress core fetches flow state data from local memory. Stores payload in DRAM. All ops in

Packet data aggregated and streamed by interface

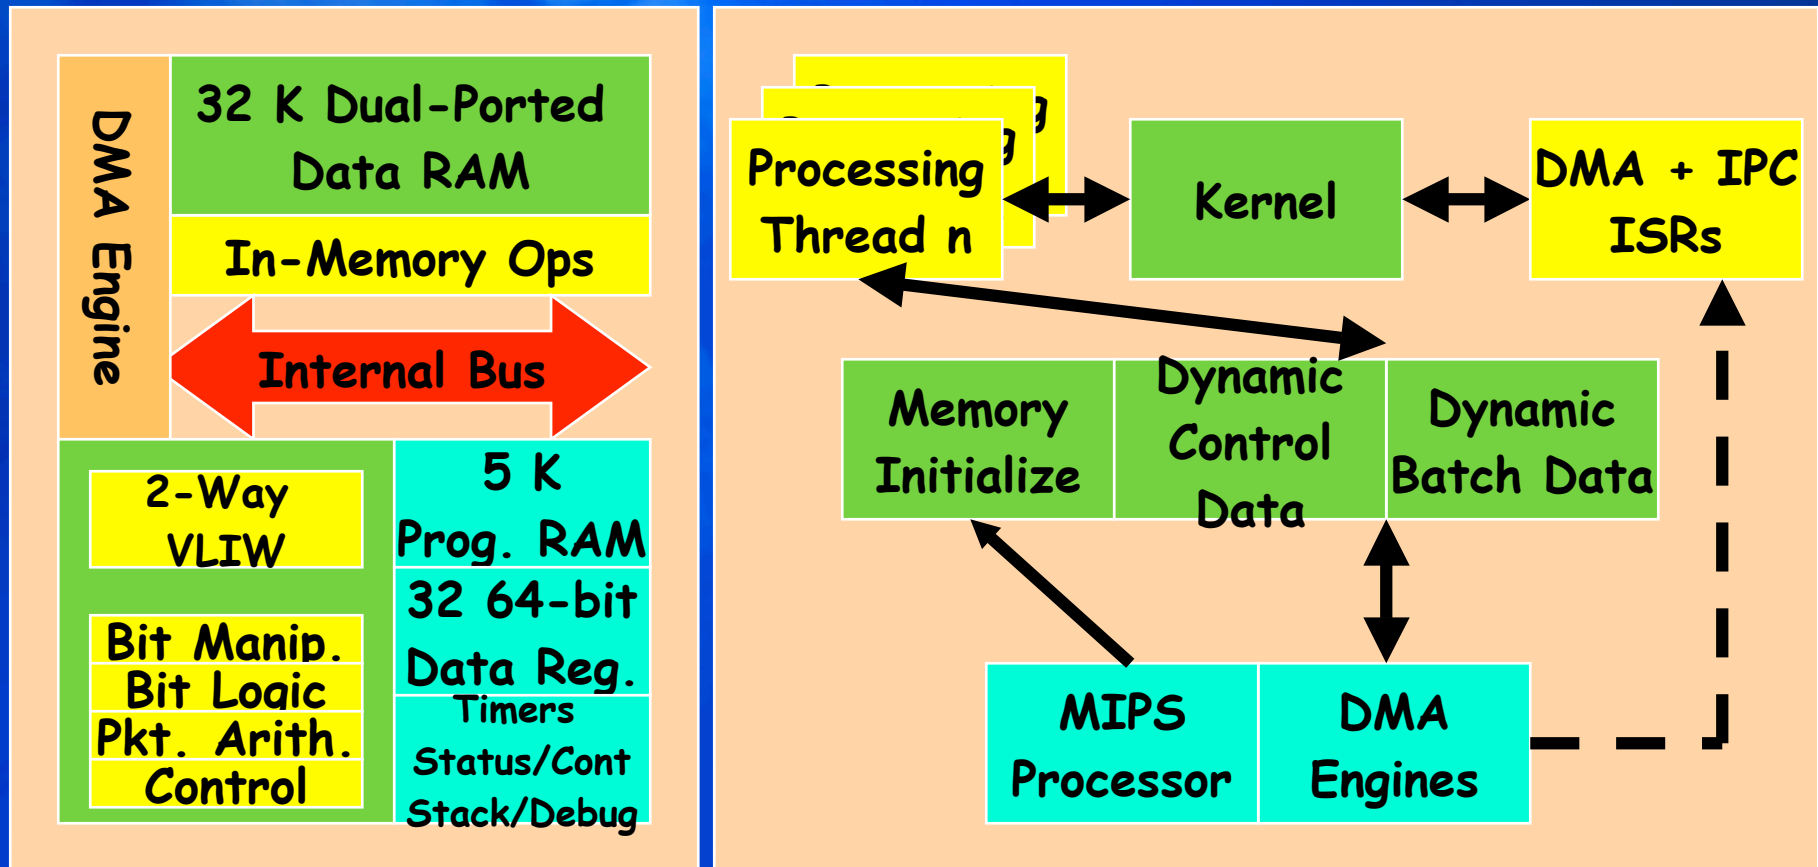
N I F

M A C

DSP Farm to Packet Interface Data Packet Flow



Packet Engine Architecture



RISC-based core with custom instructions for functional acceleration

Instruction Set Class/ Application Matrix

	Flow Assoc.	Payload Ops.	Header Ops	Shaping Control
Control Flow		X	X	X
Bit Manip.	X		X	X
Packet L & A	X	X	X	
In Memory	X	X	X	X

Each class of instructions is applicable to multiple flow processing operations

Packet Engine ISA

- **Bit Manipulation**
 - Header explosion/aggregation, Endian rearrange, Bit insert/extract
- **Packet Logic and Arithmetic Operations**
 - Checksums, CRCs, In-place addition, Reg. controlled adds, In-place pattern matching, Bit counts, Shift-logic operations
- **Memory Operations**
 - In-memory copy, Tree/List/Pointer manipulation, Multi-register loads, Byte-aligned operations
- **Program flow control**
 - Predicate-based branches, Hardware loopds, Variable-count loops, Multi-index loops, Multi-way decisions
- **Multi-instruction links**
 - Coordinate control/data instructions

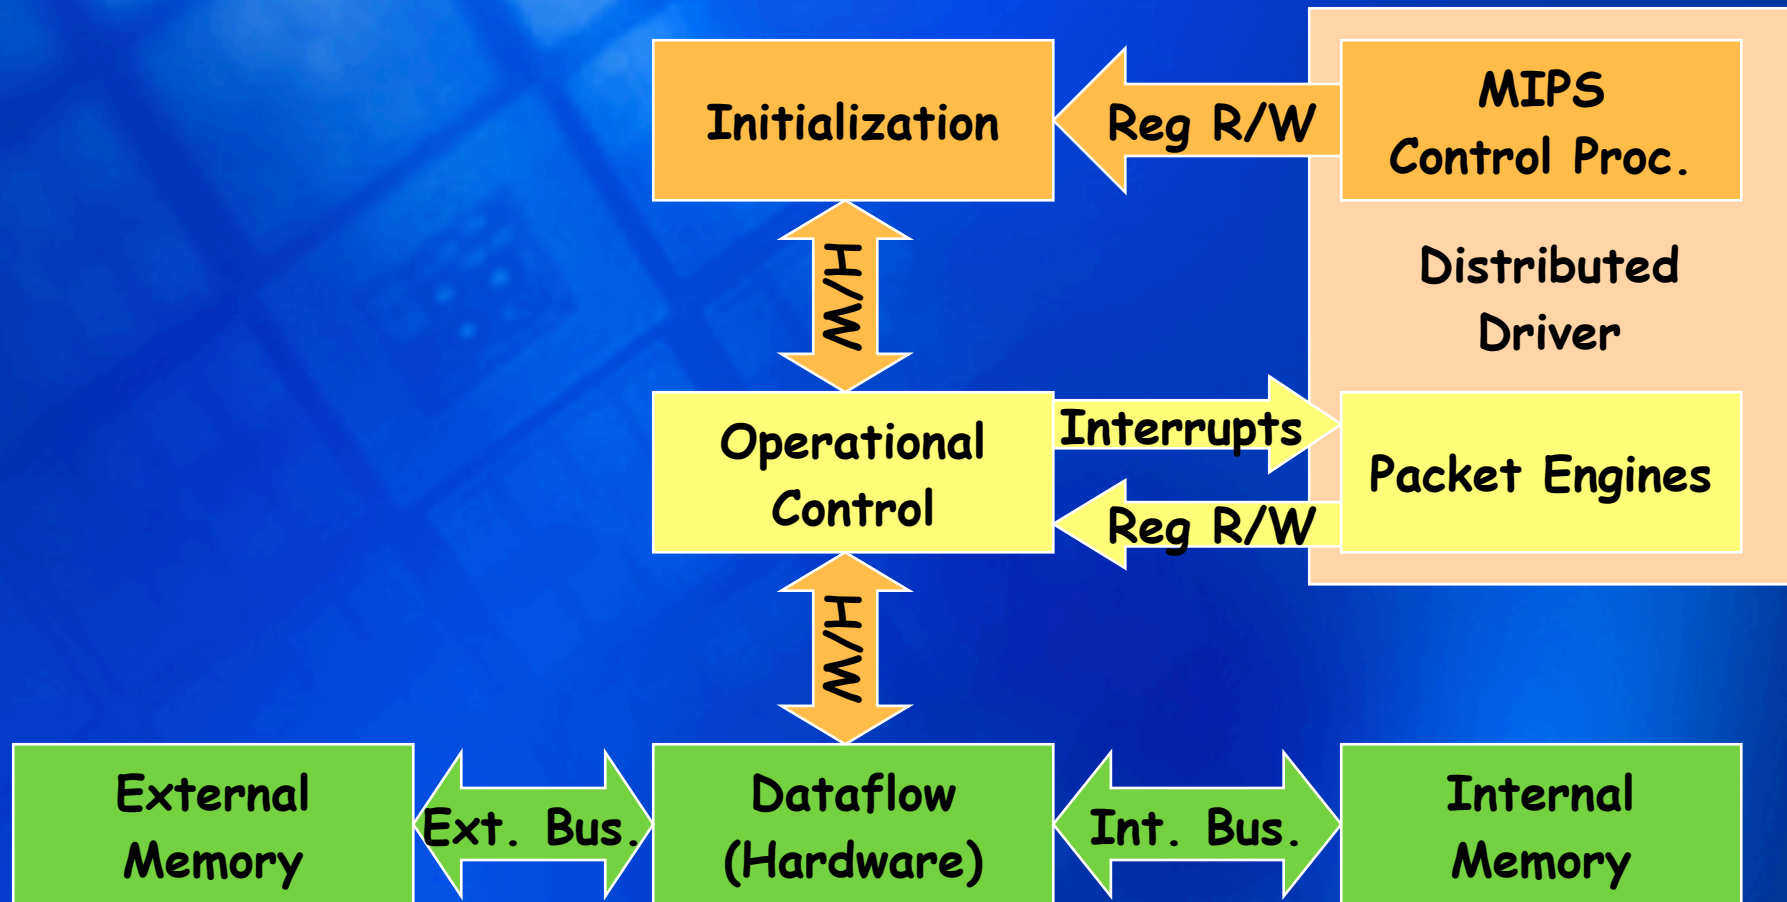
ISA Applications

- **Flow association**
 - CAM emulation, Hash functions, Heap-based look-up, Table look-up, Tree/List-based look-up
- **Payload operations**
 - Packetization/Cellification, AAL2 SAR/RTP Muxing, Payload encoding/encryption, Jitter buffer insertion
- **Header Operations**
 - RTP/IP/UDP/ATM Header parsing/assembly, Checksum/CRC calculation/verification, Field updates
- **Traffic shaping, flow control**
 - Jitter buffer extraction, Calendar queuing, Statistics updates, GCRA algorithms, Hierarchical queuing

ISA Application Examples

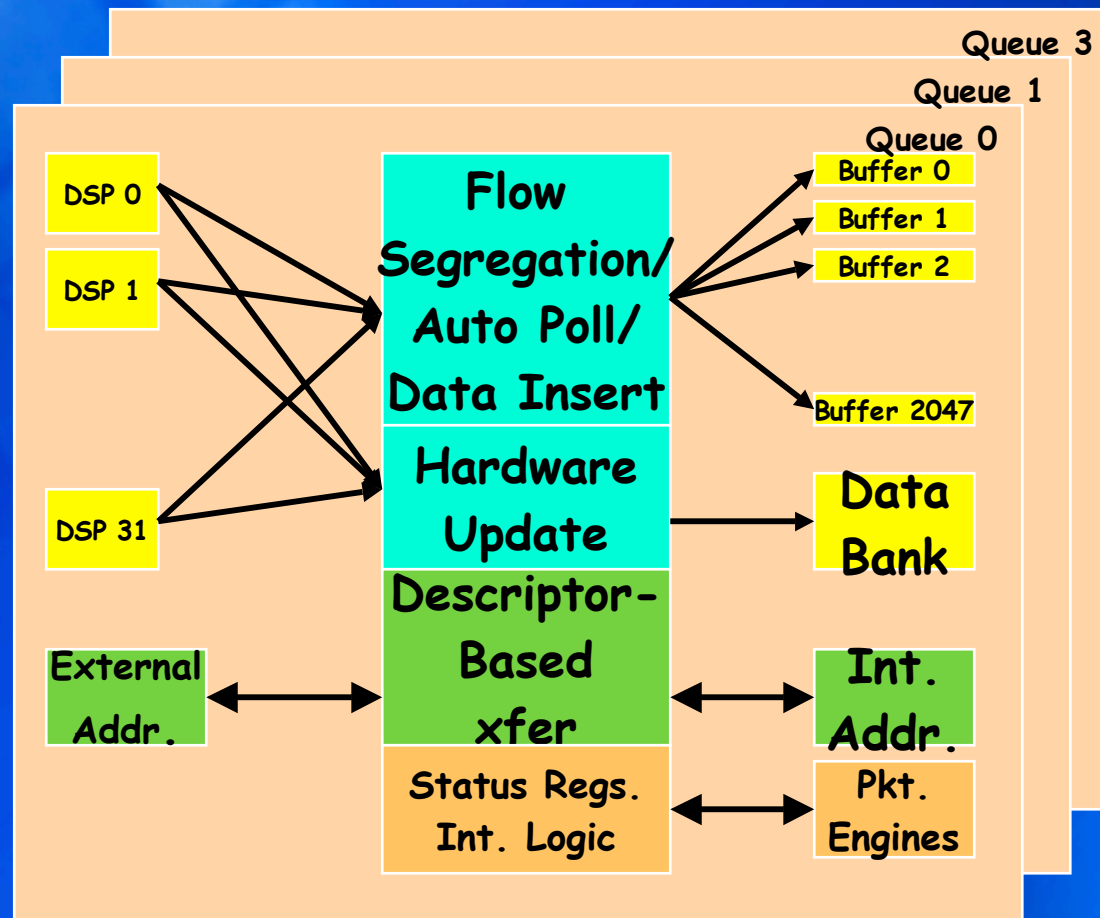
- **Checksum for IP/UDP**
 - ELOOP //Identifies number of bytes to be loaded**
 - LST EPLD // Loads data into register file**
 - SMAD //Single cycle checksum up to 16 bytes**
- **Tree traversal for flow association**
 - ELOOP //Controls tree traversal/**
 - LST ETREE //Looks for match/follows appropriate branch/NOPs**
- **Software SAR for AAL5/AAL2 (I363.2, I366.2, RFC 1483)**
 - VLOOP //Identifies segment for SAR**
 - LST EMCPY //Copies over data**
- **ATM Header extraction**
 - EEXTR //Explodes target data into up to 5 registers**

Interface DMA Structure



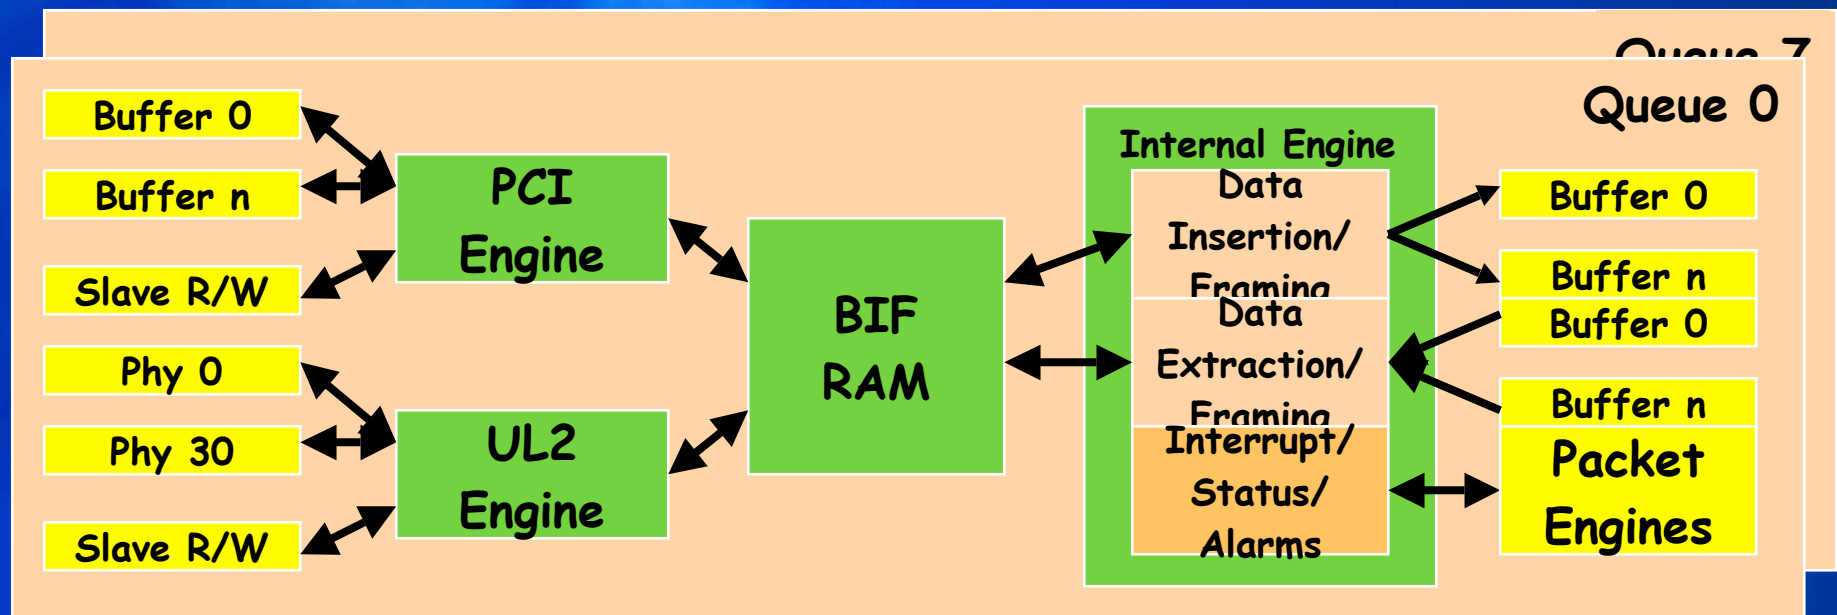
Distributed DMA engine designed to offload routine data movement tasks
Aggregates packets to reduce data movement overhead

DSP Interface Functions



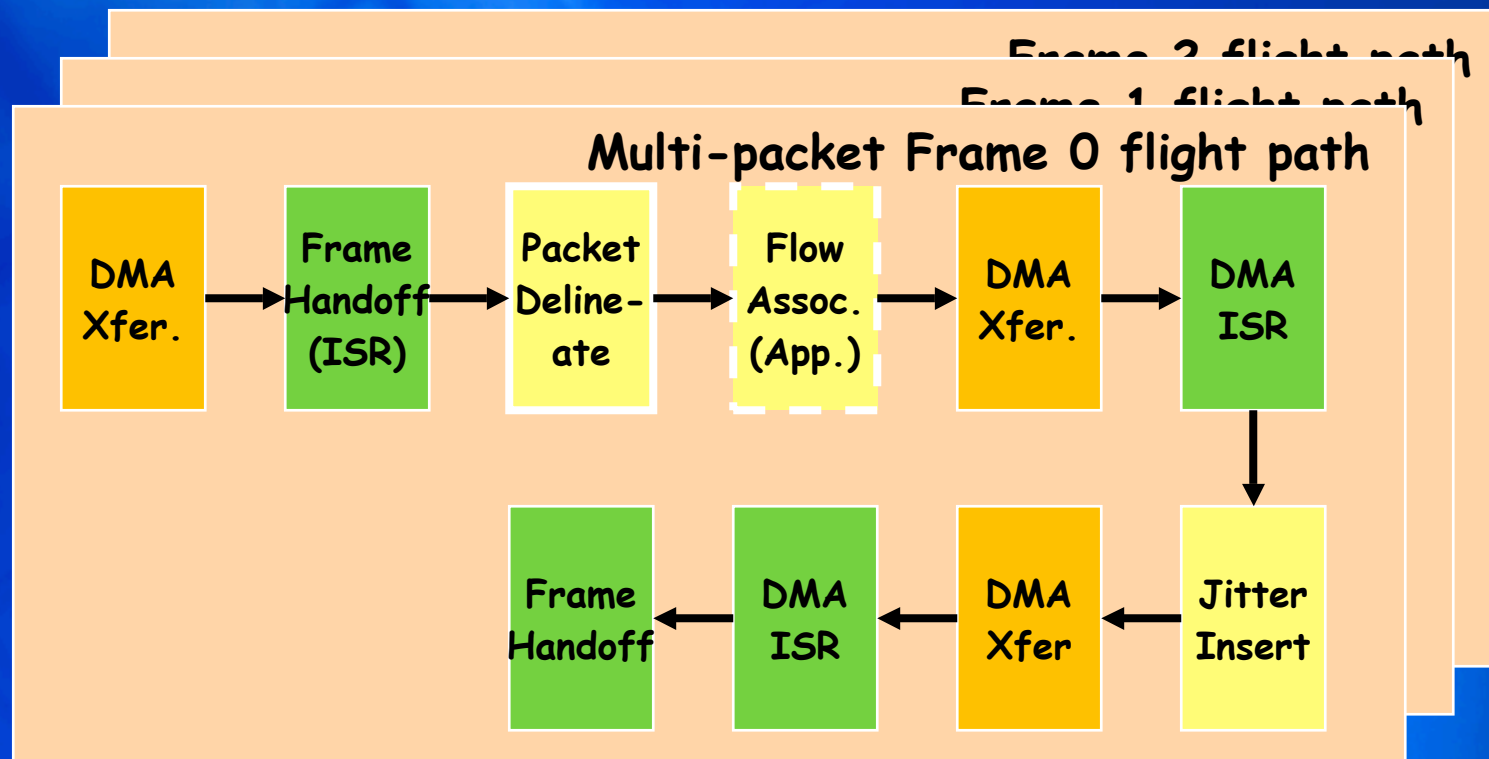
DSP interface integrates per-packet DSP push/pull and polling with QoS support

Network Interface Functions



Network interface presents multiple external interface in common format
Aggregates small packets to reduce data transport overhead
Absorbs network jitter



Packet Engine Application Architecture



Pipelined data flow enables compute/movement overlap
Memory sized to support multiple frames of data in flight in
each engine simultaneously

Magpie Physicals

Parameter	Value
Technology	0.18 μ
Packet Core Gates/Size	~1M Gates/ ~7.5mmx4.5mm
Gate/Transistor Count/Die Size	~10M/~70M/ ~15.5mmx15.5mm
Clock Rate/ Power	125MHz/ 4W/~2mW/flow
Package	625-pin PBGA



Magpie Performance

Parameter	Value
Target throughput	2016 Flows
Per flow data rate/packet rate	64 Kbps/128 Kbps/200 pps
Aggregate data/packet rate	155 Mbps/258 Mbps/403 Kpps
Cycles/packet at 125 MHz	310
Ingress actual cycles/functions	170 - flow association, header processing, replacement, checksum, flow state update
Traffic actual cycles/functions	90 – flow status check, jitter buffer timing check, packet release
Egress actual cycles/functions	130 – flow association, header replacement, checksum creation, state update, packet release

Conclusion

- **Milliflows present at high density flow aggregation nodes**
 - The bandwidth of the aggregate flow more than 1000 times that of component flows
- **Systems to service aggregation nodes are required to meet stringent constraints**
 - Real-time service for several thousand component flows under tight power constraints
- **Developed an experimental SoC IC targeting real-time milliflow aggregation**
 - Functional acceleration for primary functions in VoIP packet processing
 - Results in an effective power-performance trade-off
- **Thanks to**
 - Sunil Chaudhari, Kumar Ganapathy, Saleem Mohammadali, Jonathan Liu, Saurin Shah, David Gilbert, T.H. Lu, Sameer Nanavati, Jennifer Donnelly, Carl Alberola, Manoj Mehta