

High Performance Communication

- Ada Gavrilovska – ada@cc
- class web page
- 8803 (4803) – the ‘history’ of 8803HPC
- grading:
 - homework assignments: 15%; paper presentation: 15%;
class participation: 10%;
 - class project: 40% - *implementation*, 10% - *report*, 10% -
presentation (total of 60%)
- letter grade, P/F, audit

Material

- Lecture material:
 - no required textbook
 - most papers/documents – online
 - + handouts in class
- Project material:
 - recommended
 - [IXP2400/2800 Programming](#), *The Complete Microengine Coding Guide*, by Erik J. Johnson and Aaron Kunze
 - IXP technical documents – in /net/hp31/ixpdev
 - includes programming reference manual, hardware description, tools, etc.; also code from book examples
 - Infiniband documentation, OpenIB/OpenFabric
 - openfabrics.org
 - /usr/ofed

Administrative

- announcements:
 - online; emergency once – will send email to class
- get access to ILAB cluster
 - should be automatic for everyone in class, but send email to help@cc if not
 - will (need to) have a home dir on /net/hp31/
 - communication – ixp@cc
- get access to IHPCL
 - infiniband clusters: Polynesia and Rohan, /usr/ofed/
 - communication - ihpcl-lab@cc
- request for class mailing list pending, will post online

Course components

- Lectures – list tentative and open for suggestions

Course components

- Lectures – list tentative and open for suggestions
- Paper presentations

Course components

- Lectures – list tentative and open for suggestions
- Paper presentations
- Warm-up exercise – 2, based on IXP and IB

Course components

- Lectures – list tentative and open for suggestions
- Paper presentations
- Warm-up exercise – 2, based on IXP and IB
- Project
 - define by Sep. 20th, may work alone or in pairs
 - midterm report – 1page status update, after Midterm Recess – Oct. 18th
 - last two/three classes – project presentations (Dec. 4th, 6th)
 - final project report and demo due by Wed. Dec. 13th (reserved period for final exam)

Project info

- possible platforms: IXP (recommended), Infiniband (OpenIB/OpenFabric), kernel hacking for multi-core platforms... other – FPGA-based, programmable routers,
- “sponsored by industry partners”
- Ideas
 - online list will be updated, talk to me, look at class papers for inspiration...
 - what do you want to do?
 - related to current work, or want to experiment with potential research ideas

Scope

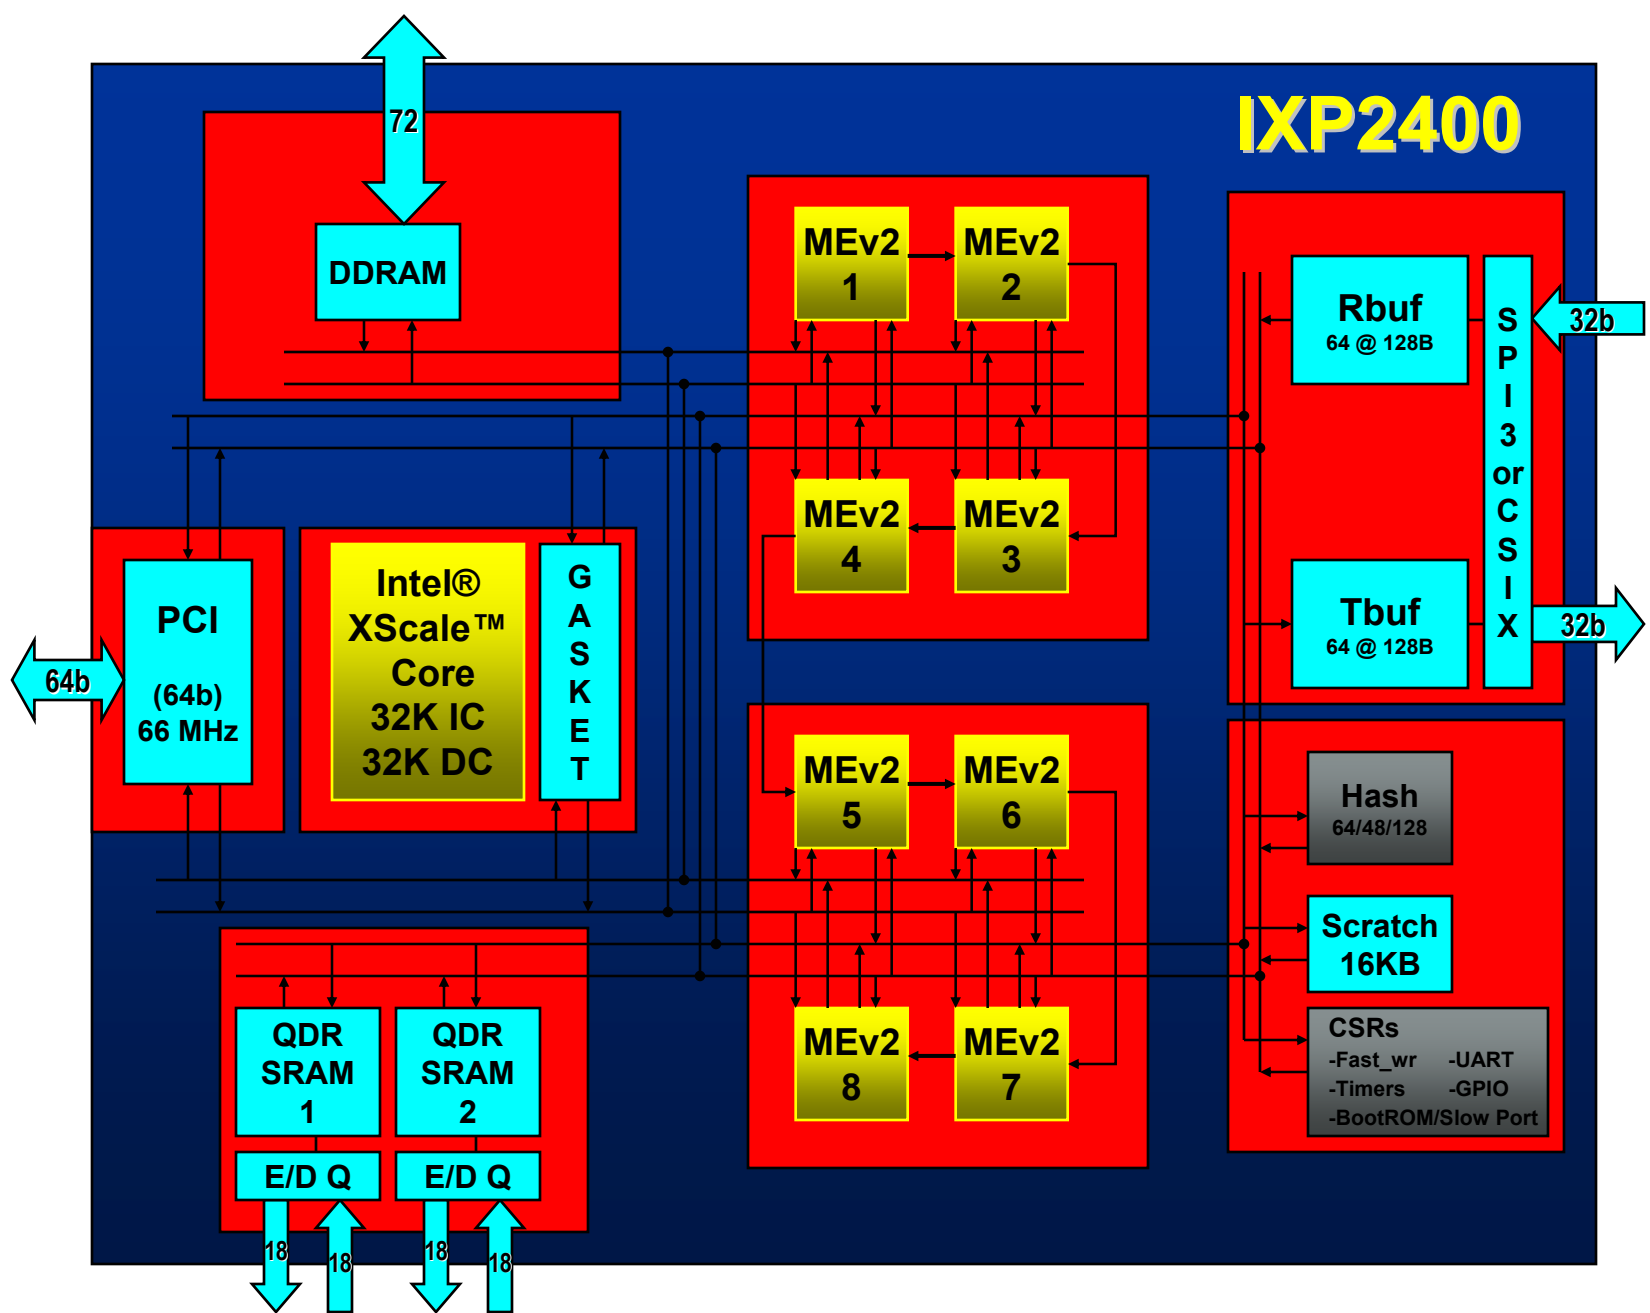
- Application domain
 - scientific – traditional parallel applications, simulations, remote collaborations
 - enterprise – commercial systems – banks, airlines, dynamic/personalized delivery of remote content...
 - dynamic – inputs/outputs, interactions, needs/interests, policies and regulations
- High performance communication for high performance applications
 - should translate to high availability, reliability, scalability, consistency, manageability, security...
 - should enable high rate of operations/transactions

- communication
 - getting the needed data from sources to destinations
- high-performance
 - high data-rate, low-delay
 - quality => latency and bandwidth, jitter and error rates;
 - but also quality in terms of application-specific metrics => data and processing needed now
- need mechanisms at all levels

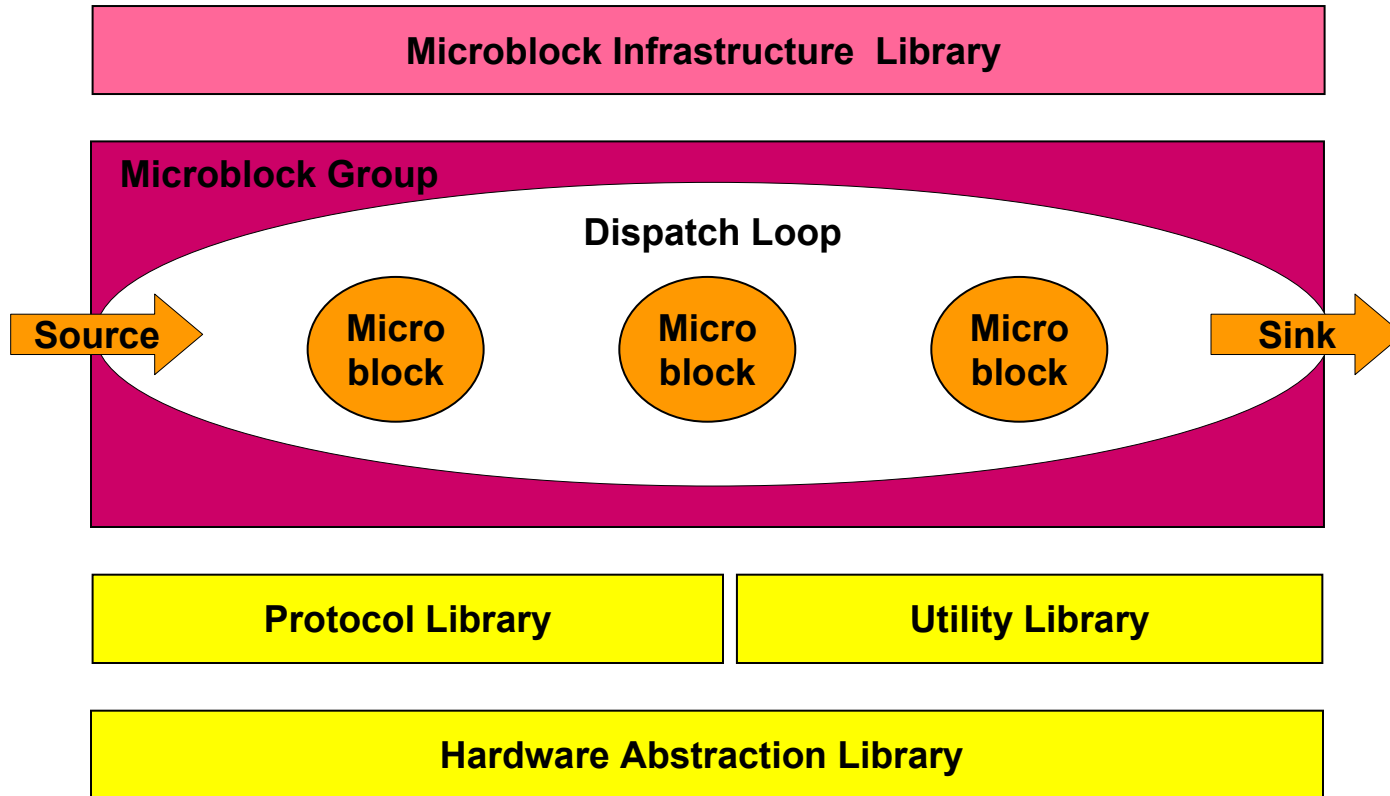
Lecture schedule

- Intro & overview of tools
- War of the Interconnects
- Programmable Networking vs. NetASICs
- IP and Multiprocessor Networking Protocols
- High Performance System Services/Mechanisms for Communications
- HP Communication Services on Multi-core platforms
- Programming Models & Environments
- Middleware and Applications
- Evaluation and Modeling Tools

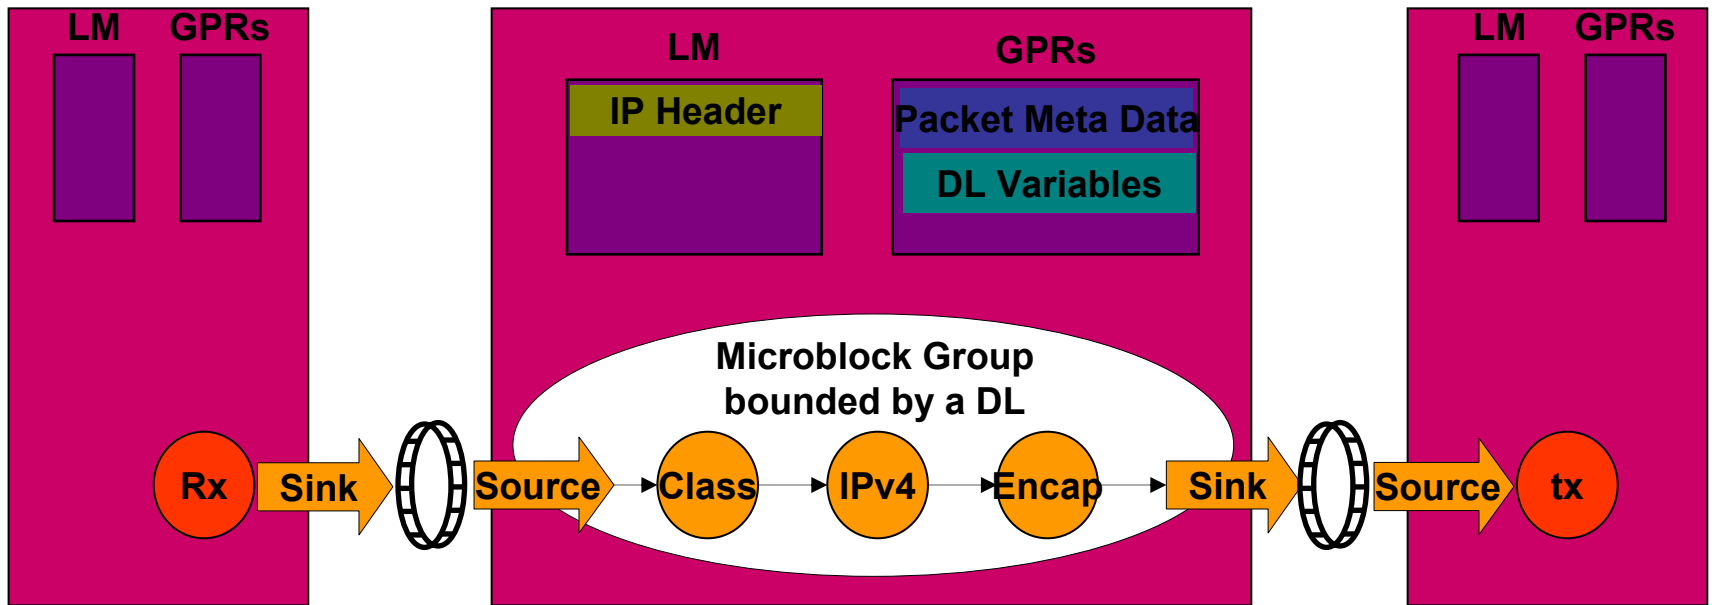
IXP2400



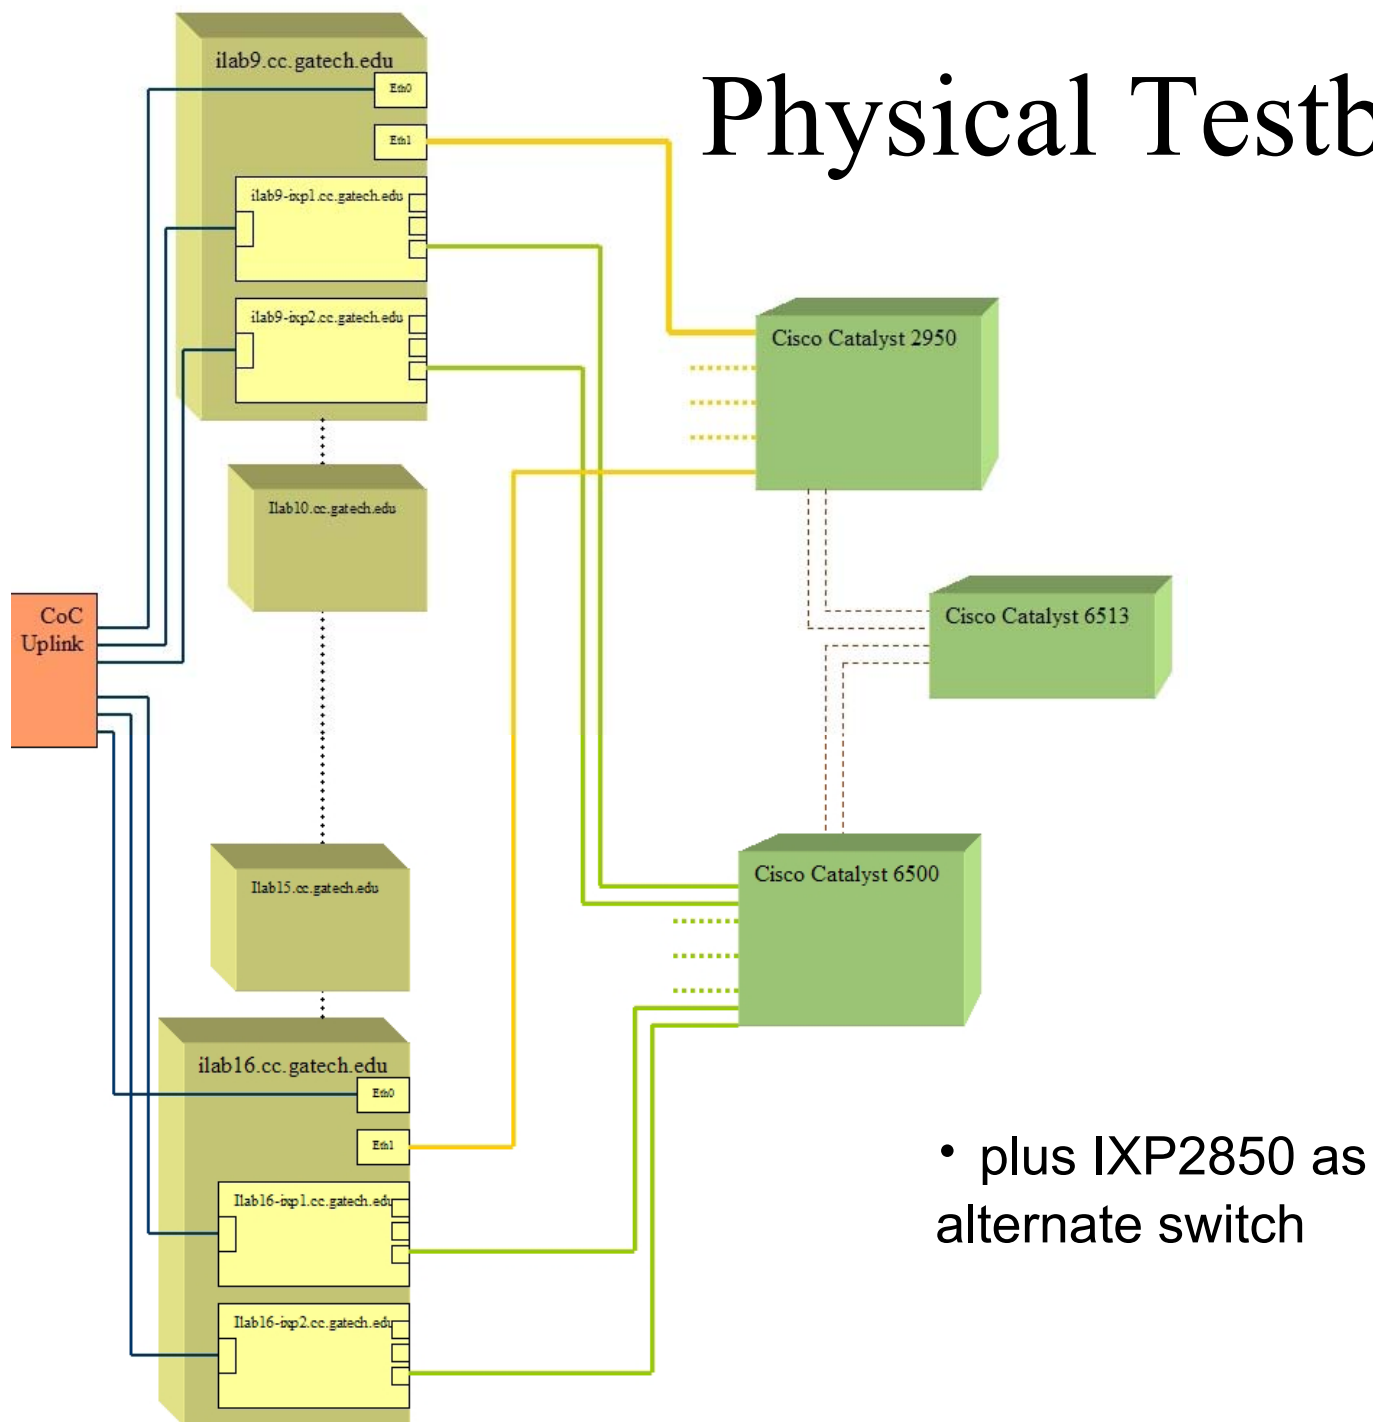
Microengine Programming Model



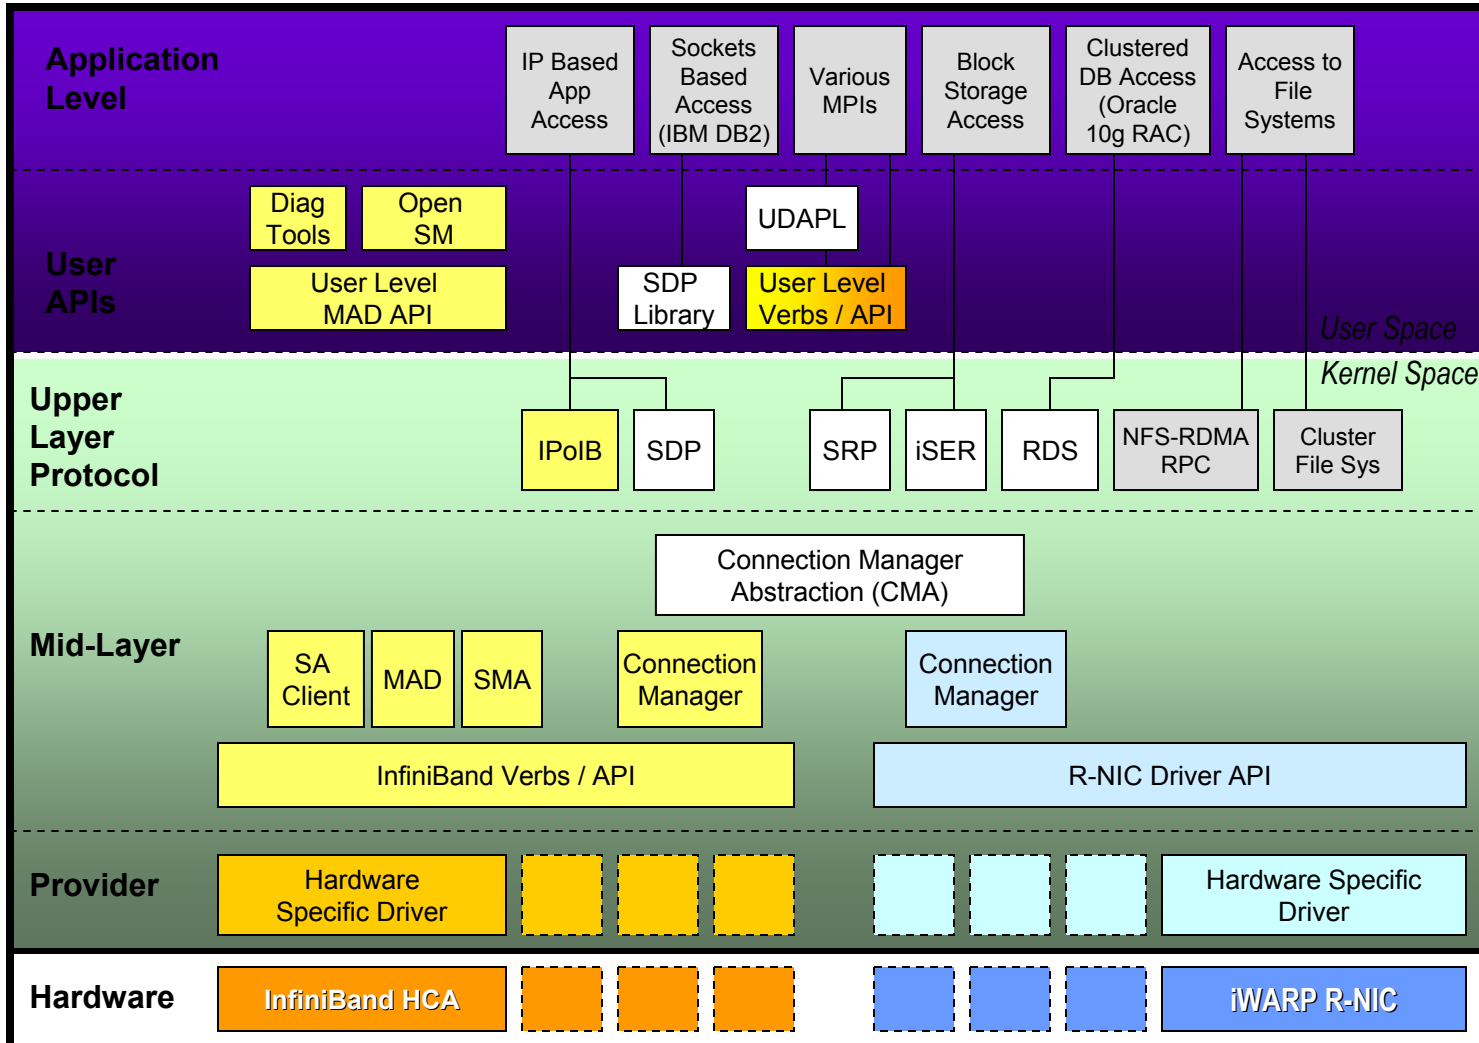
Microblock Architecture



Physical Testbed



OpenFabrics Software Stack



SA	Subnet Administrator
MAD	Management Datagram
SMA	Subnet Manager Agent
PMA	Performance Manager Agent
IPoIB	IP over InfiniBand
SDP	Sockets Direct Protocol
SRP	SCSI RDMA Protocol (Initiator)
iSER	iSCSI RDMA Protocol (Initiator)
RDS	Reliable Datagram Service
UDAPL	User Direct Access Programming Lib
HCA	Host Channel Adapter
R-NIC	RDMA NIC

Key	Common	Apps & Access Methods for using OF Stack
	InfiniBand	
	iWARP	