

KU EECS 800

Survivable, Resilient, and Disruption Tolerant Networking

James P.G. Sterbenz

The University of Kansas Department of Electrical Engineering and Computer Science

jpgs@eecs.ku.edu

3036 Eaton

++1 785 864 8846

<http://www.ittc.ku.edu/~jpgs/courses/eecs800>

Survivability

Definition

Survivability is the capability of a system to fulfill its mission in a timely manner, even in the presence of *attacks* or failures [CMU SEI]

- Assignment:
readings and references for other definitions

Fault Tolerance

Scope

- Fault models do not hold under malicious attack
 - we cannot assume independence and random failures
- Examples:
 - 1988 Hinsdale Illinois Bell central office fire
 - 100K customers lose service for *weeks*
 - also major disruptions in long distance, 800, 911, cellular, ATC
 - various DNS and BGP meltdowns
 - 2003 NE power failure (SCADA supervisory control and data acquisition)
 - 2005 SBB power outage
- *FT* necessary but *not* sufficient for survivability

Fault Tolerance

Examples

- 1988 Hinsdale Illinois Bell central office fire
 - 100K customers lose service for *weeks*
 - also major disruptions in
 - long distance
 - 800
 - 911
 - cellular
 - ATC for O'Hare
- [...]

Fault Tolerance

Examples

- 2003 Northeast US power failure
 - SCADA: supervisory control and data acquisition
- [...]

Fault Tolerance

Examples

- 2005 SBB/CFF/FFS power failure
 - SBB (Schweizere Bundesbahn) power failure
- [...]

Disruption Tolerance

Definition

Disruption tolerance

is the ability for end-to-end applications to operate even when network connectivity is not strong (weak, episodic, or asymmetric) and the network is unable to provide stable end-to-end paths [JPGS]

Resilience

Definition

Resilience is the capability of the network maintain an acceptable level of service in the face of various challenges to normal operation, *including legitimate traffic* [DH+JPGS]

- Resilience is a superset of survivability capabilities
 - therefore, survivability is necessary but may not be sufficient

Challenges to Normal Operation

Resilience \supset Survivability/DT \supset FT

- Unusual (legitimate) traffic load
 - e.g. flash crowds
- High-mobility of nodes and subnets
- Wireless channels
 - weak, episodic, asymmetric connectivity
- Long delay paths
 - distance, transmission delay
- Attacks against
 - network hardware, software, protocol infrastructure
- Large-scale natural disasters
- Misconfiguration and operational errors
- Natural faults of network components



Survivability & Disruption Tolerance

Motivation and Threats

- Network & applications should remain operational
 - when the network is under attack
 - in challenged environments: mobile / wireless / long-delay

Survivability & Disruption Tolerance

Motivation and Threats

- Network & applications should remain operational
 - when the network is under attack
 - in challenged environments: mobile / wireless / long-delay
 - particularly critical and lifeline services
 - we can't depend on the Internet for this today
- Attacks against the physical infrastructure
 - natural disasters
e.g. hurricanes, earthquakes, ice storms, tsunami, floods
 - targetted attacks (terrorism or warfare)
- Attacks against protocol and software infrastructure
 - recreational crackers: denial of service
 - industrial espionage and sabotage
 - cyber-terrorism and information warfare

Resilience

Motivation and Threats

- Network & applications should remain operational
 - when the network is under attack
 - in challenged environments
 - under abnormal traffic load
 - DDOS
 - flash crowds (e.g. slashdot, breaking news event)
- Distinguishing attacks can be *very* hard
 - a sufficiently sophisticated distributed denial of service attack is indistinguishable from legitimate traffic
 - it doesn't matter
 - to other users (cross traffic)
 - network provider (resource exhaustion)

Resilience and Survivability/DT

Motivating Scenarios

- Mountain rescue [Lancaster 6NET]
- Mobile library [Lancaster 6NET]
- Flash crowd and DDOS attack [Lancaster TA]
- Snowmobile nomads [Sweden]
- Ubiquitous personal computing & communication
- Tactical networks for homeland security and warfare
- Resilient backbone network infrastructure

Resilient Network Scenarios

Mountain Rescue

- Mountain rescue [Lancaster 6NET]
 - Cockermouth Mountain Rescue Team, English Lake District
 - rescue teams share mobile vehicle
 - one or more team per rescue
 - team uses vehicle as base station
 - episodically connected with base and one-another
 - may need to multihop to base station
 - base station network (weak or episodic connectivity)
 - wireless relay to Internet
 - wireless connection to one-another

Resilient Network Scenarios

Mobile Library

- Mobile library [Lancaster 6NET]
 - hi-tech bookmobile
 - travels to villages with weak or no connectivity
 - electronic as well as printed books and other content
 - prefetched and cached
 - preloaded
 - pre-reserved by customer
 - anticipated by librarian
 - Internet access
 - prefetched, cached, and scheduled preloaded content
 - normal access when connectivity strong enough

Resilient Network Scenarios

Flash Crowd and DDOS Attack

- Flash crowd / DDOS attack [Lancaster]
 - unexpected volume of traffic targetting a server
 - sophisticated DDOS undetectable
 - indistinguishable from legitimate traffic
 - DDOS detection only raises bar of attack difficulty
 - effects are the same as flash crowd
 - DOS against server
 - damage to cross traffic
 - exhaustion of network resources
 - set of programmable-network mechanisms to tolerate *both*
 - traffic engineering
 - route and path manipulation
 - transport layer intervention (e.g. SYN dropping)

Communication Environment


Impact of Wireless Channel

- Open channel subject to *attack*
 - eavesdropping
 - network and traffic analysis
 - interference
 - jamming and denial of service
 - injection of bogus signalling and control messages
- Weak, intermittent, and episodic connectivity
 - limited bandwidth of shared medium
 - time-varying available bandwidth
 - noise, weather (latter for free-space laser as well as RF)
 - episodic connectivity
 - channel fades between bit errors & failed links in consequence
 - difficult to achieve routing convergence

Communication Environment

Impact of Mobility


- Dynamic nodes and topologies
 - changing links, clustering, and federation topology
 - difficult to achieve routing convergence
- Control loop delay
 - mobility may exceed ability of control loops to react
- QOS



Communication Environment

Impact of Mobility

- Dynamic nodes and topologies
 - changing links, clustering, and federation topology
 - difficult to achieve routing convergence
- Control loop delay
 - mobility may exceed ability of control loops to react
- Impacts QOS
 - changes in inter-node distance
 - requires power adaptation
 - changes density and impacts degree of connectivity
 - latency issues (routing optimisations temporary)



Communication Environment

Impact of Mobility

- Dynamic nodes and topologies
 - changing links, clustering, and federation topology
 - difficult to achieve routing convergence
- Control loop delay
 - mobility may exceed ability of control loops to react
- Impacts QOS
 - changes in inter-node distance
 - requires power adaptation
 - changes density and impacts degree of connectivity
 - latency issues (routing optimisations temporary)



3

Communication Environment

Impact of High Latency

- Long inter-application delay appears to be disruption
 - long path (c)
 - store-and-forward queueing due to episodic connectivity
 - latency masking techniques mitigate: caching, prefetching
 - but *don't always* help
- Severely impacts transport and network protocols
 - signalling latencies dominate at high data rates
 - very long control loops
 - long delays may cause data transfer to stall (window-based)
 - wrapped sequence number spaces
 - high-bandwidth- \times -delay products
 - real-time reaction to many bits in flight difficult or impossible
 - massive buffering required for error control

Resilience and Survivability

Assumptions and Challenges₁

- Problem cannot be solved at physical and link layers
 - assume that best physical, MAC, link techniques in use
 - diminishing returns on further research
- Strong connectivity will not always be achievable
 - economics and policy preclude connectivity everywhere
 - faraday cages for security
 - caves
 - nomadic hunters in northern Sweden; search and rescue in UK
- Network / security infrastructure may be unavailable
 - node failure or overrun (capture)
 - radio silence or jammed channel (enemy, cracker, DDOS)
 - compromised node software

Resilience and Survivability

Assumptions and Challenges₂

- Very long delay inevitable in some scenarios
 - path (speed of light) latency
 - satellite links
 - interplanetary (and intergalactic?) Internet
 - object transmission delay
 - large objects over modest data rates
 - weakly connected and congested links
 - store-and-forward over episodically connected paths
- Security and survivability are not binary choices
 - level of security must be traded against resource cost ...
 - limited node power
 - limited channel bandwidth
 - ... based on application requirements and user desires

Resilience and Survivability

Requirements and Goals

- Resilience and survivability goals for *network*
- Resilience goals for *services*
- Disruption tolerance goals for *applications*

Resilience and Survivability

Requirements and Goals: Network

- Resilience and survivability goals for *network*
 - resistance to attack and traffic anomalies
 - recognition when attack/anomaly has occurred
 - recovery from attack/anomaly after occurrence
 - refinement in future response to attack/anomaly
- Metrics

Resilience and Survivability

Requirements and Goals: Services

- Resilience goals for *services*
 - accessible with graceful performance degradation
 - correct operation
 - resistance from attack
 - recovery from performance degradation
- Metrics

Resilience and Survivability

Requirements and Goals: Applications

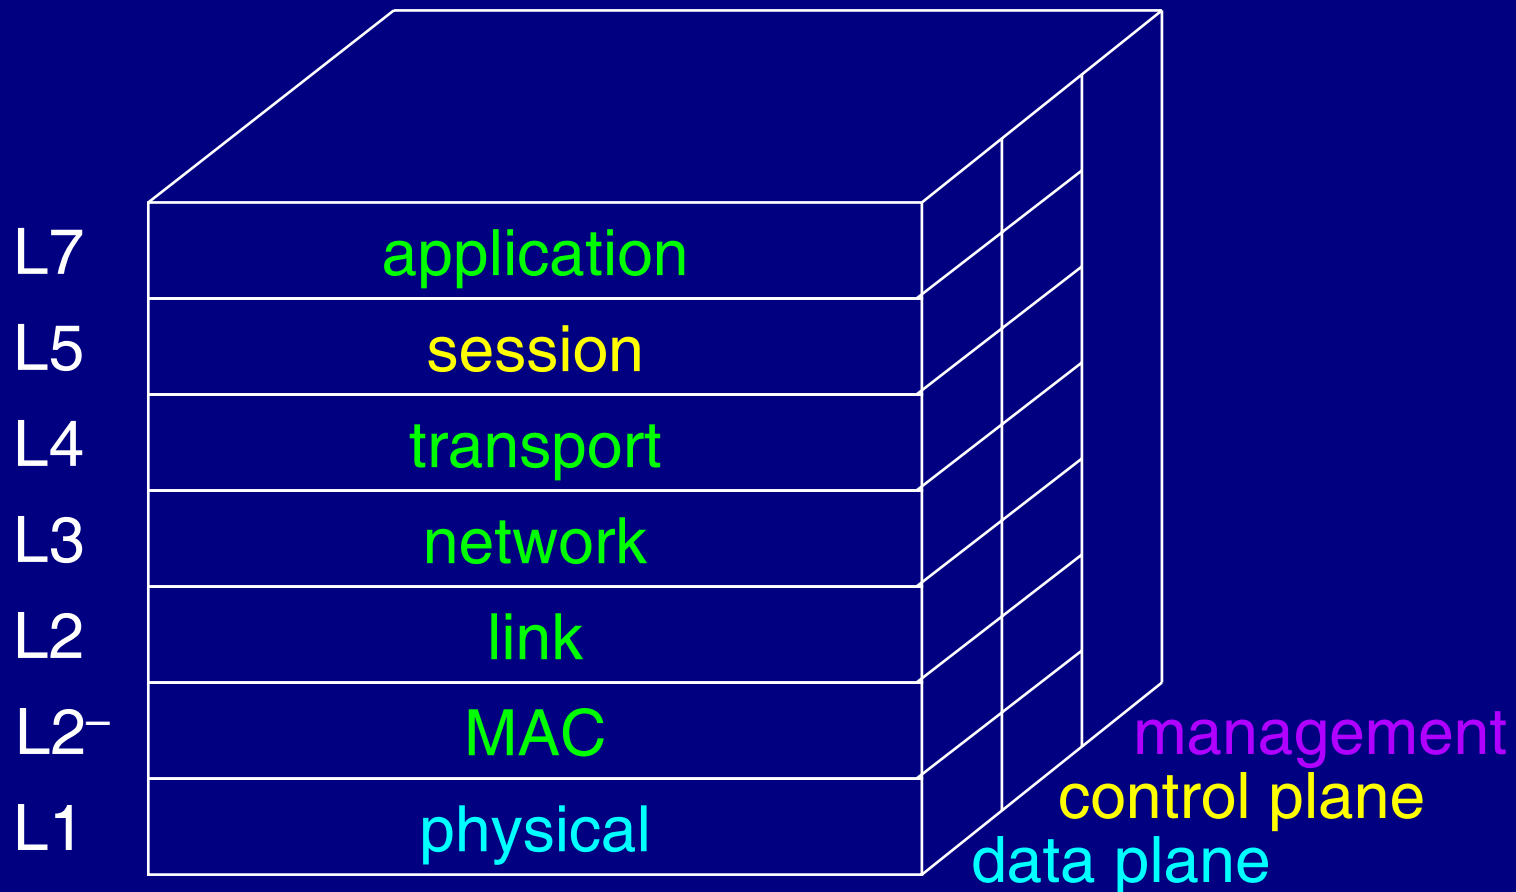
- Disruption tolerance goals for *applications*
 - access to information by the user or application
 - e.g. Web browsing
 - maintenance of end-to-end communication association
 - e.g. video- or teleconference
 - distributed processing and networked storage
- Metrics

Protocol Layering Planes

- Data
 - information transfer
- Control
 - signalling to control information transfer, including:
 - flow or connection establishment/modification/termination
 - error control
 - flow and congestion control
 - correspond to data layers
- Management
 - monitoring and management of network and its elements
 - cuts across all layers


Protocol Layering

Hybrid Layer/Plane Cube




Link Layer

Layer/Plane Cube Model



Link Layer


Link Definition



- *Link* is the interconnection between *nodes*
 - *intermediate systems* (switches or routers)
 - *end systems* (or hosts)

Link Layer

Link Protocol



- Link protocol
 - is responsible for per hop transfer of data *frame*
 - assume dedicated links for now: no MAC [see lecture M]

Link Layer

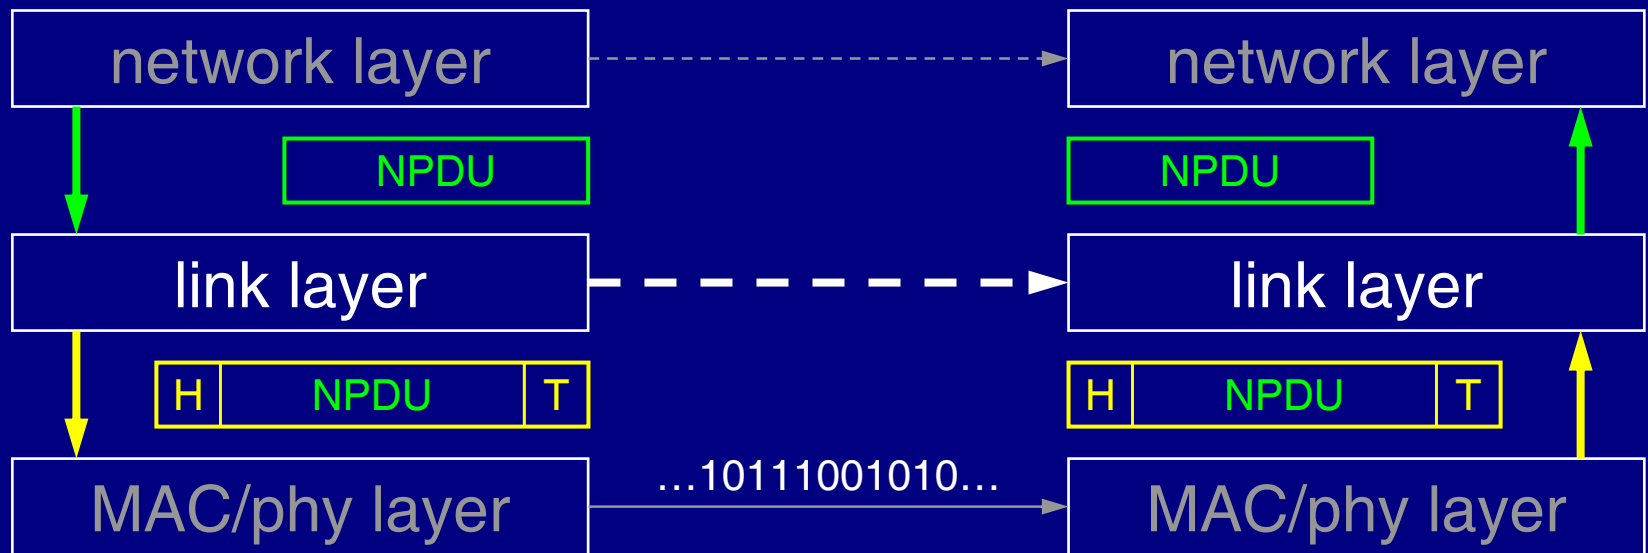
Service and Interfaces

- Link layer is HBH analog of E2E transport layer
 - transport layer (L4) transfers packets E2E
- Link layer (L2) service to network layer (L3)
 - transfer frame HBH (hop-by-hop)
 - sender: encapsulate packet into frame and transmit
 - receiver: receive frame and decapsulate into packet
 - error checking / optional correction or retransmission
 - recall end-to-end arguments:
E2E reliability with HBH error control only for performance
 - flow control possible but not generally needed at link layer
 - parameter negotiation typical (e.g. data rate)
- Link layer multiplexing and switching

Link Layer

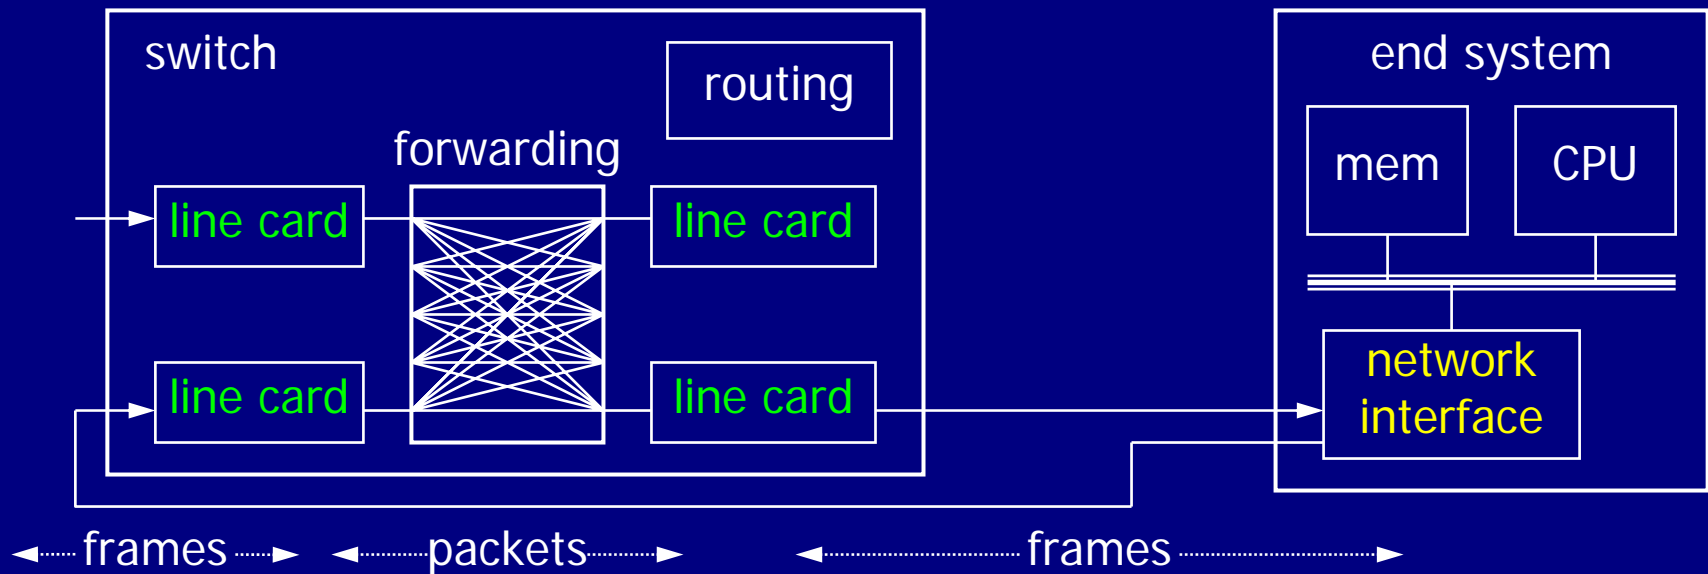
Service and Interfaces

- Link layer *frame* encapsulates network layer *packet*
 - packet (NPDU – network layer protocol data unit)
 - frame (LPDU link layer protocol data unit)
 - frame = header + packet + (trailer)



Link Layer

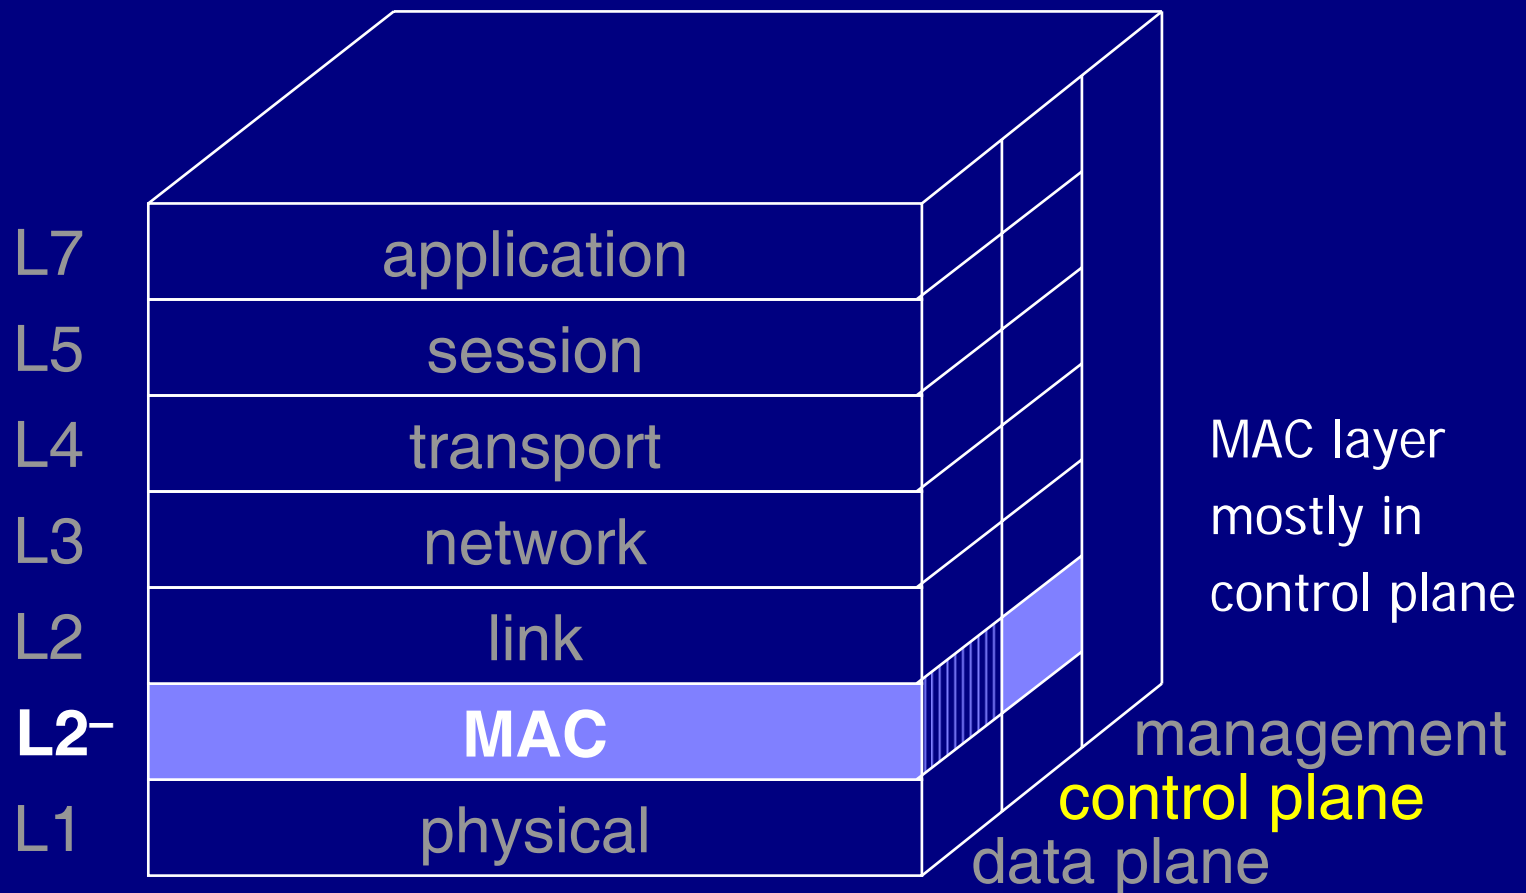
Functional Placement



- Link layer functionality per line interface
 - end system: **network interface (NIC)**
 - switch/router: **line card**

Medium Access Layer

Layer/Plane Cube Model



Medium Access Layer

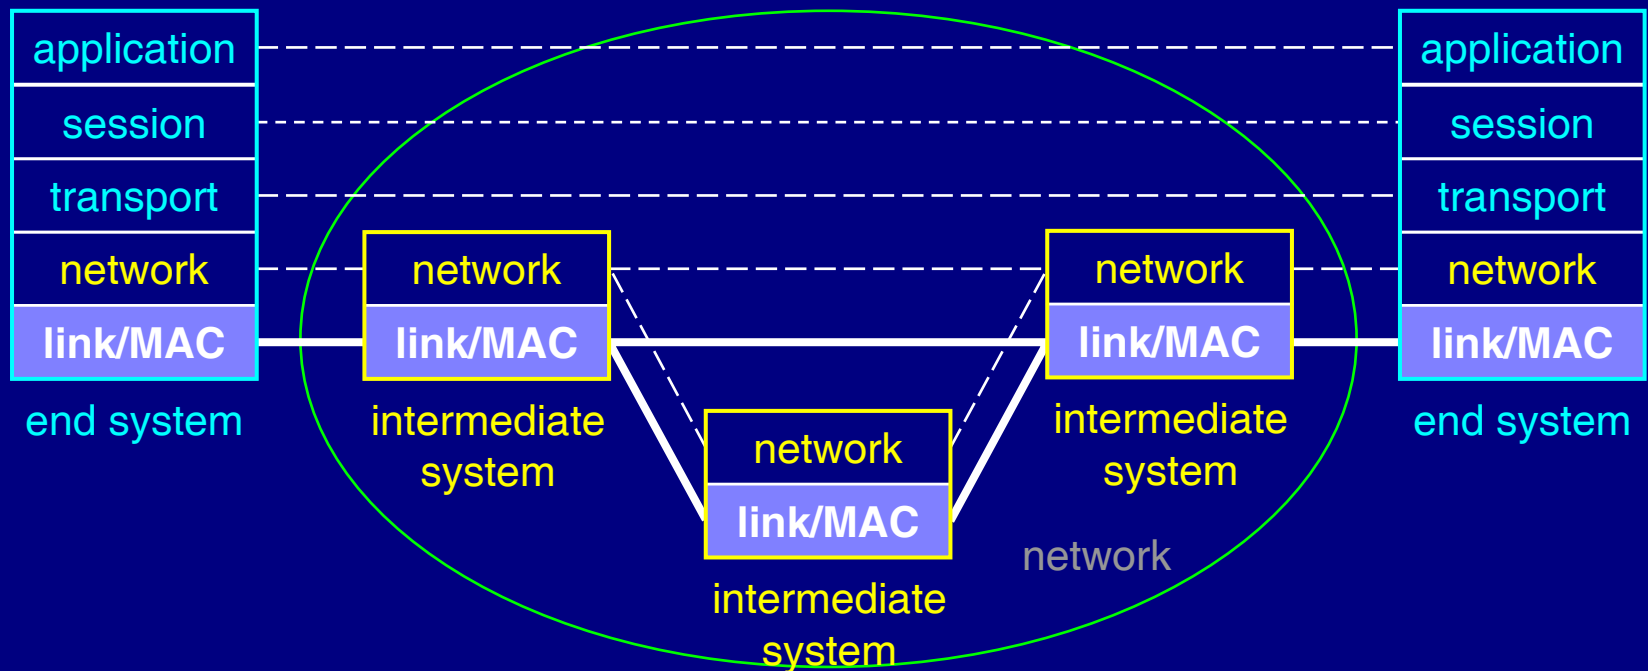
Link Definition



- *Medium access control* arbitrates a *channel* in shared medium (free space, guided wire) or fiber among stations

MAC Layer

MAC Protocol



- MAC protocol (or algorithm)
 - is responsible for determining when node can transmit *frame*
 - may fully distributed or coordinated

MAC Layer


Service and Interfaces

- MAC is in-between physical layer 1 and link layer 2
 - lower link sublayer in IEEE 802 model
- MAC layer is mostly in control plane
 - control over when to transmit a L2 frame
 - but may have its own encapsulation (e.g. IEEE 802 legacy)
- Mac layer service to link layer (L2)
 - MAC layer encapsulate/decapsulate if appropriate
 - initiate transfer of frame into the medium

Link Layer

Service and Interfaces

- MAC layer *frame* may encapsulate link layer *frame*
 - done for link layer / MAC protocol independence
 - IEEE 802: 802.2 LLC (logical link control) over 802.x MAC



MAC Algorithms

Major Types

- Channel partitioning
 - TDMA, FDMA, WDMA
 - best under high deterministic load
- Random access
 - ALOHA, CSMA
 - best under light load
- Coordinated access (taking turns)
 - polling, token passing
 - combines benefits of both partitioning and random access
- Spread spectrum
 - CDMA

Assignment for Next Class

- Read [F2003]
- Become familiar with SIGCOMM W-DTN papers
 - we'll discuss W-DTN next class
- Follow references in [SK+2002] and [KS+2003]
- Begin researching scenarios and past failures
 - e.g. NE power failure, 9/11, SBB, Hinsdale CO, *others*

End of Foils