- CS 470
  - Multiprocessing
    - OpenMP
      - Pragmas
        - parallel
        - for
        - section(s)
        - task
        - single
        - master
        - critical
        - barrier
      - Clauses
        - default/private/shared
        - reduction
        - schedule
        - collapse
        - firstprivate/lastprivate
        - nowait
      - Functions
        - omp_get_wtime
        - omp_get_num_threads
        - omp_get_max_threads
      - Loop-carried dependencies
      - Tasks and teams
      - Locks
    - Parallel languages
      - Productivity vs. performance
      - Partitioned global address spaces (PGAS)
      - High-Performance Fortran
      - CAF and UPC
      - X10 and Fortress
      - Chapel
      - Python and Julia
    - Explicit vs. implicit
    - Threads vs. processes
    - Thread safety
      - Non-determinism
      - Race conditions
      - Deadlock
    - Synchronization
      - Mutual exclusion
      - Locks
      - Semaphores
      - Conditions
      - Monitors
      - Barriers
  - Parallel patterns
    - Task vs. data decomposition
    - Shared-memory vs. distributed-memory
    - Locality
      - Spatial vs. temporal
      - NUMA effects
      - Caching
      - Mirroring
      - Content delivery networks
      - Space-filling curves
    - Foster's methodology
      - Partitioning
      - Communication
      - Aggregation
      - Mapping
    - Communication patterns
      - Naturally ("embarrassingly") parallel
      - Reduction trees
      - Nearest-neighbor
      - Producer/consumer
      - Map/reduce
      - Pipelines and streams
    - Collective operations
      - Broadcast
      - Reduction
      - Scatter
      - Gather
      - Allgather
      - Allreduce
      - All-to-all
    - Matrix operations
      - Sparse vs. dense
      - Linear system solvers
      - Linear algebra
  - Performance analysis
    - Speedup and efficiency
      - Amdahl's and Gustafson's laws
      - Linear speedup
      - Critical path analysis
      - Temporal vs. spatial locality
      - Weak vs. strong scaling
      - CPU time vs. wall time
    - Communication overhead
      - Bandwidth
      - Latency
      - Bisection
      - Contention
    - Energy usage
      - Energy (work) vs. power (rate)
      - Power caps
      - Dynamic voltage frequency scaling
    - Analysis tools
      - Debuggers
      - Profilers
        - Hardware counters
        - Timer resolution
        - Sampling
        - Overhead
        - Perturbation
        - Skid
      - Tool frameworks
      - Performance modeling
      - Roofline model
      - Autotuning
  - System architectures
    - Flynn's taxonomy
      - SISD
      - SIMD
      - MIMD
      - SPMD
    - Instruction-level parallelism
      - von Neumann bottleneck
      - Pipelining instructions
      - Superscalar processing
      - Speculative execution
      - Vector processing
    - Shared memory
      - Pthreads
      - Java threads
      - Windows threads
      - OpenMP
    - Coprocessors and accelerators
      - GPUs / GPGPUs
      - CUDA
      - OpenACC
    - Manycore
      - Intel Phi
      - Sunway
    - Distributed clusters
      - OpenMPI and MPICH
      - Homogeneous vs. heterogeneous
      - Hybrid w/ accelerators
      - Topologies
        - Bus
        - Crossbar switches
        - Star
        - Ring
        - Grid / Mesh
        - Torus
        - Hypercube
        - Fat trees
      - Interconnects
        - Ethernet
        - InfiniBand
        - OmniPath
      - Supercomputers
    - Wide-area networks
      - End-to-end principle
      - Sockets
      - OSI model
      - QoS concerns
      - Routing
        - Circuit switching vs. packet switching
        - Unicast
        - Multicast
        - Broadcast
      - Web protocols
        - IP / DNS
        - TCP / UDP
        - HTTP / HTML
        - SSL / TLS
        - NTP
        - XML / SOAP / JSON
      - Peer-to-peer
        - Bittorrent
        - Tor
        - Freenet
    - Clouds
      - Infrastructure-as-a-service
      - Virtualization
        - Type-1 vs. type-2 hypervisors
        - Virtual machines
        - Containers / Docker
      - Amazon AWS
    - Grids
      - Work-sharing
      - Condor
      - GIMPS
    - Novel architectures
      - Memory-centric
      - Neuromorphic
      - Quantum
      - Optical
  - Distributed issues
    - Naming
      - Flat namespaces
      - Hierarchical namespaces
      - IPv4, IPv6, and DNS
      - Partitioned global address spaces (PGAS)
      - Overlay networks
      - Distributed hash tables
        - Virtual address space
        - Finger / lookup tables
        - Chord
    - Synchronization
      - Message passing
      - Clocks
        - Physical
        - Lamport clocks
        - Vector clocks
        - NTP
      - Barriers
      - Consensus protocols
        - Transactions
        - Elections
        - One-phase vs. two-phase
        - Paxos
    - Replication and consistency
      - Partial vs. total orderings
      - Data-centric
        - Continuous
        - Sequential
        - Causal
      - Client-centric
        - Monotonic reads
        - Monotonic writes
        - Read-your-writes
        - Writes-follow-reads
      - Distributed version control
    - Fault tolerance
      - CAP theorem
        - Consistency
          - Strong
          - Eventual
          - Weak
        - Availability
          - Active / passive
          - Active / active
        - Partition tolerance
      - Soft vs. hard failure
      - Permanent vs. intermittent vs. transient faults
      - MTBF and FIT
      - Failure types
        - Crash
        - Omission
        - Timing
        - Response
        - Arbitrary (Byzantine)
      - Failure handling
        - Detection
        - Prevention
        - Avoidance
        - Recovery
        - Techniques
          - DMR vs. TMR
          - Checksums / hashes
          - Hamming codes
          - Reed-Solomon codes
          - Checkpointing
    - Security
      - Attacks
        - Brute force password cracking
        - Replay attacks
        - Man-in-the-middle attacks
      - Principle of least privilege
        - Trust
        - Policies
      - Encryption
        - One-way hash functions
        - Cryptographic systems
        - Symmetric vs. asymmetric
        - MD5 / SHA
        - DES / RSA
      - Authentication
        - Shared-key challenge/response
        - Needham-Schroeder
        - Kerberos
        - Key exchange parties
        - Diffie-Helman key exchange
        - Certificate authorities
      - Authorization
        - Firewalls
        - Access control lists
        - LDAP and AD
      - Auditing
        - Append-only logs
        - Blockchains (Bitcoin)
    - File systems
      - Design issues
        - File-level vs. block-level
        - Remote access vs. upload/download
        - Centralized vs. decentralized
        - Symmetric vs. asymmetric
        - Striping
      - Remote procedure calls
        - Function stubs
        - Parameter marshalling
        - Synchronous vs. asynchronous
      - Networked file systems
        - Exports
        - Mounts
        - Static vs. automatic
      - Protocols
        - NFS
        - AFS
        - GoogleFS
        - Lustre
        - Bittorrent
        - Freenet
    - Middleware
      - Scheduling
        - SLURM
      - Monitoring
      - Load balancing
      - Checkpoint/restart