Beyond the Pale:
The Role of Abstraction in the Internet

M. Steenstrup
steenie@rcn.com

The Essential Tension

User objectives and constraints:

- communicate successfully with intended recipients independent of time, place, and content
- expend minimal effort and cost to obtain desired service

Provider objectives and constraints:

- commercial: maximize profits by providing sufficient services to satisfy customers with minimal expenditure at highest price tolerated
- community: offer affordable interconnectivity among and Internet access to users, but may have limited resources
- ad hoc: share personal communications resources with others with expectation of tit-for-tat

Dimensions of Communication

Focus:

- who: communication with specific users or locations
- what: communication of specific information or for specific functions

Sources and destinations:

- point-to-point
- one-to-many, many-to-many, many-to-one

Services:

- delay sensitivity
- loss tolerance
- throughput

Size, Heterogeneity, and Dynamics

Traffic flows:

- sources and destinations
- duration, rate, and burstiness
- service needs and provider preferences

Communications devices:

- intended purpose and ownership
- computation, memory, and communications technologies
- location and trajectory
- activity schedule and available power

Communications channels:

- transmission media
- accessibility and price of bandwidth
- existence, error, and loss characteristics

The Reality

Communications resources:

- may be severely limited in certain places at certain times depending on capacity of communications equipment, current demand, environmental interactions, and enforcement of usage restrictions

Network state and control information:

- large quantity
- subject to frequent and potentially unpredictable changes
- control likely based on delayed, incomplete, and inaccurate state information

Types of communication:

- wide variety
- desire to support all of it in one (inter)network

Architectural Implications for Network Control

Resource management:

- over-provisioning may be impractical or even impossible
- distribution and replication of control functions among devices to reduce response delay and increase fault tolerance and according to device capabilities

Information containment:

- reduce amount of state and control information transmitted, processed, and stored in the network

Robust control:

- degrade gracefully with magnitude of difference between actual and observed state
- move network toward desired performance most of the time
- global optimization is likely to be impractical and impossible

Reducing the Frequency of State Updates

Set lower bound on time interval between successively reported changes in state

Check periodically for perceived changes in local state

Set lower bound on magnitude of change that triggers generation of a state update

Filter samples of local state (e.g., weighted averaging according to recency of samples) to reduce volatility of perceived state

Reducing the Number of Transmissions per Update

Set upper bound on number of hops traversed by a state update

Establish bounded region outside of which a particular state update does not propagate

Designate covering set of nodes to relay a particular state update

Relay a state update contingent on its expected utility to recipient nodes

Compressing State Information

Minimize number of bits to represent state information

Abstract state concerning multiple entities to yield state for a single aggregate entity

Hierarchical Abstraction for Networking

Benefits:

- reduction in the amount of state and control information that is transmitted, processed, and stored in the network
- availability of state and control information at multiple levels of granularity with level of detail dependent on proximity

Costs:

- control decisions based on abstracted information may differ from those based on detailed information
- may result in inefficient use of network resources and inability to find appropriate resources for traffic flows when such resources exist

General Case for Routing

Aggregation of communications devices and channels:

- different sets of devices may be organized by different aggregation algorithms
- which devices are aggregated depend on the tolerated size of an aggregate and interconnectivity, similarity of capabilities and usage policies, and ownership of the devices
- a device derives its address from the aggregation hierarchy, and different devices may reside at different depths

General Case for Routing

Abstraction of aggregate state:

- each aggregate may be subject to a different state abstraction algorithm and the devices and channels that connect adjacent aggregates are themselves abstracted
- what state information is available to a device depends on its location in the aggregation hierarchy and the state abstractions and information-hiding policies of the aggregates

General Case for Routing

Distribution of control functions over devices:

- which device performs which functions depends upon its capabilities, its location with respect to the boundaries of the aggregates of which it is a member, the transmission costs of distributing and requesting state and control information required to perform the functions, and the frequency with which the functions are expected to be performed
- redundancy of control functions across devices depends on the desired responsiveness

A Brief History

1970s:

- focus: reduce the size of a node's packet-forwarding table
- N reduced to c $\log_c N$ for c aggregates per parent aggregate and N nodes in the network

1980s:

- mobile packet-radio networks, ARPANET areas, Internet autonomous systems
- focus: reduce the amount of routing information transmitted throughout the network

1990s:

- mobile packet-radio networks and internetworks
- focus: routing as constraint satisfaction with multilevel abstraction of state for device and channel aggregates

In Practice

Limited deployment:

- Internet autonomous systems, OSPF areas
- one level of abstraction

Reasons:

- unclear whether there is a need for the most general case
- perceived as complicated to implement and manage
- service provider and equipment vendor inertia
- we tend to respond to existing rather than anticipated problems
- confusion between hierarchy resulting from aggregation of devices and channels and abstraction of state and hierarchy determined by geographical coverage of networks
- poor marketing

Network Topology

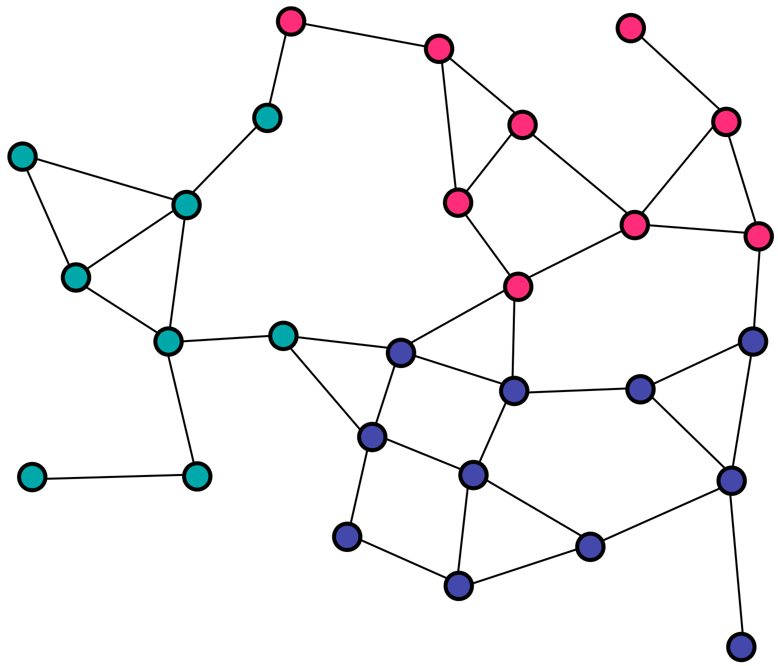

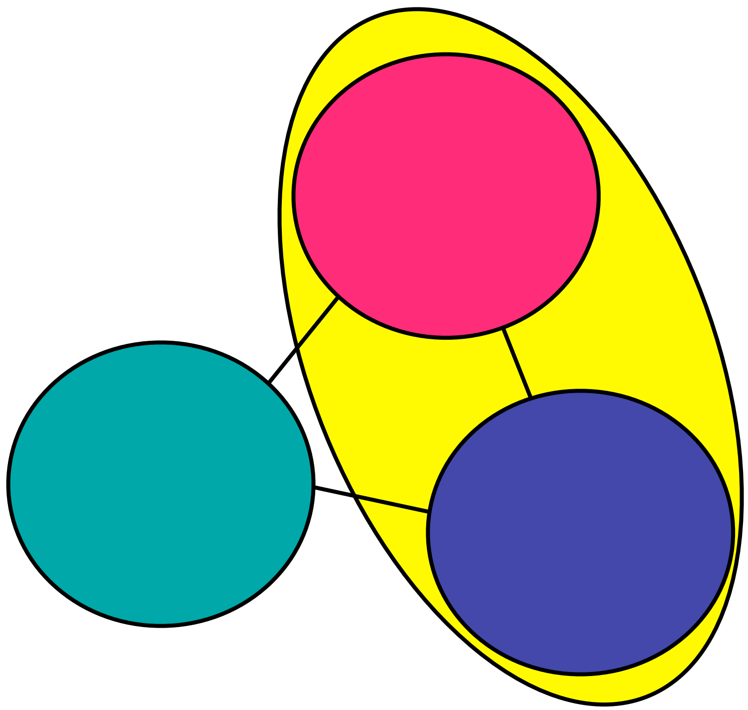

detailed view

abstracted view

Network Topology

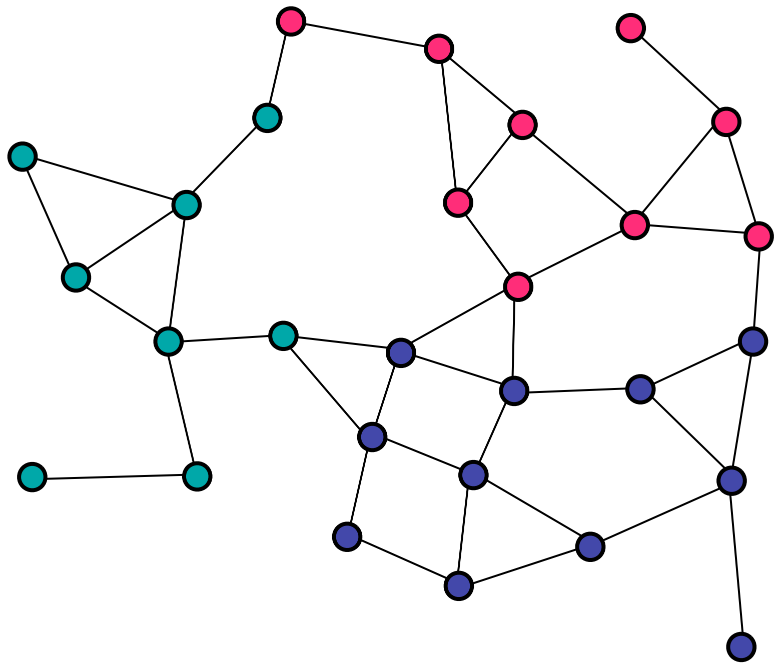

detailed view

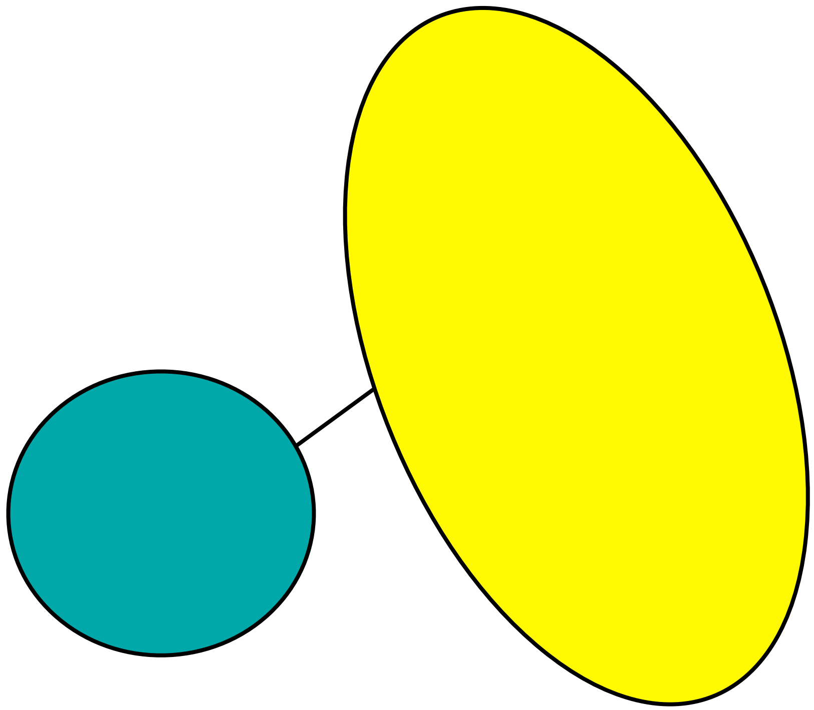

further abstracted view

Individual Device's View of Topology

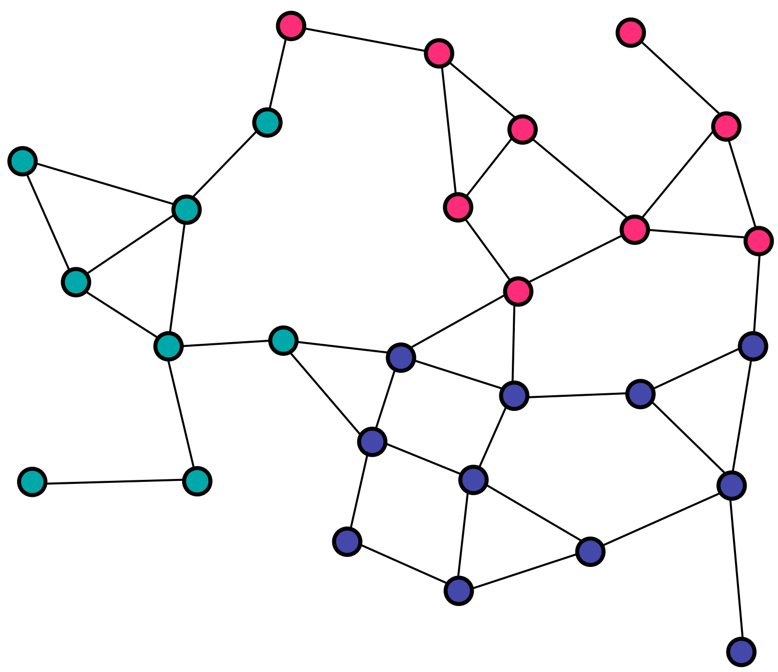

without abstraction

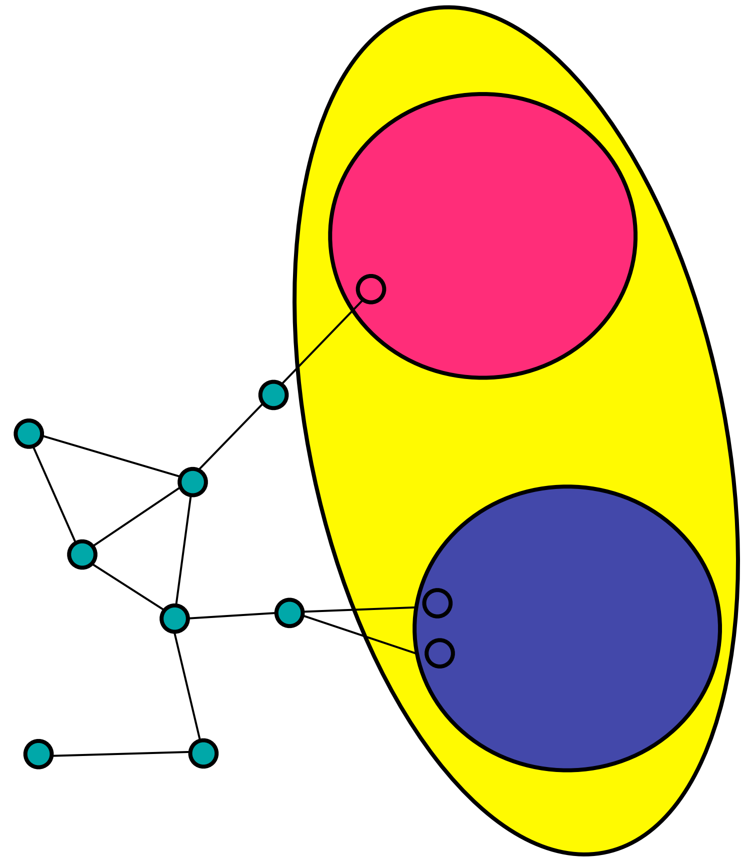

with abstraction

Costs of Abstraction

Quantity of routing information at each node:

- O(N) with routing information about each node distributed to all N nodes
- $O(c \log_c N)$ with routing information about each aggregate distributed to all nodes in the same parent aggregate, c child aggregates per parent
- O(c), with routing information about each aggregate distributed to representatives of each aggregate in the same parent aggregate

Quality of routes:

- worst case: route length may increase by orders of magnitude
- usually, increase in route length is at most a couple of hops

Types of Hierarchical Routing

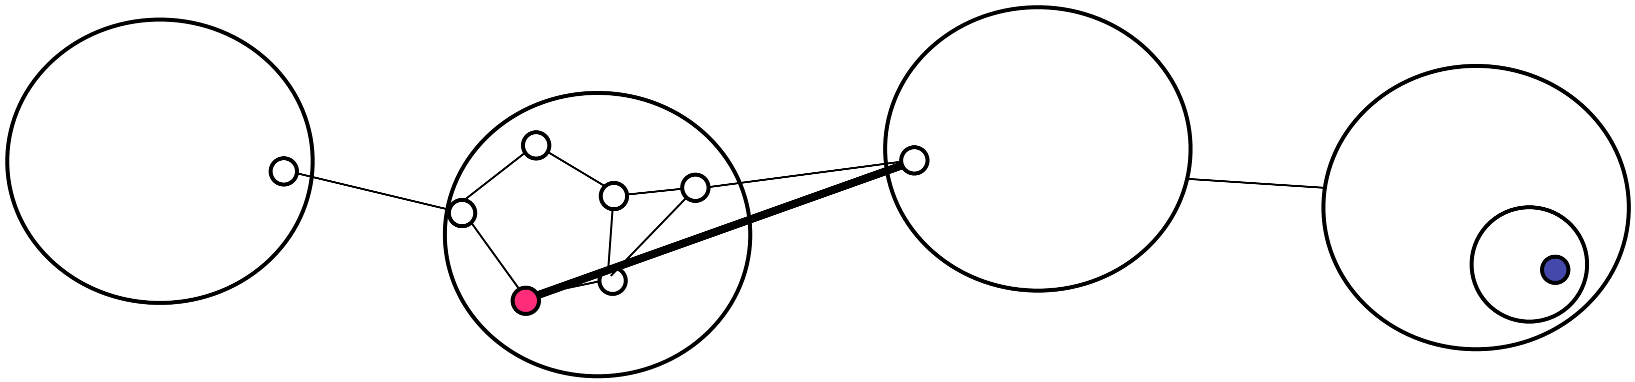

strict-hierarchical

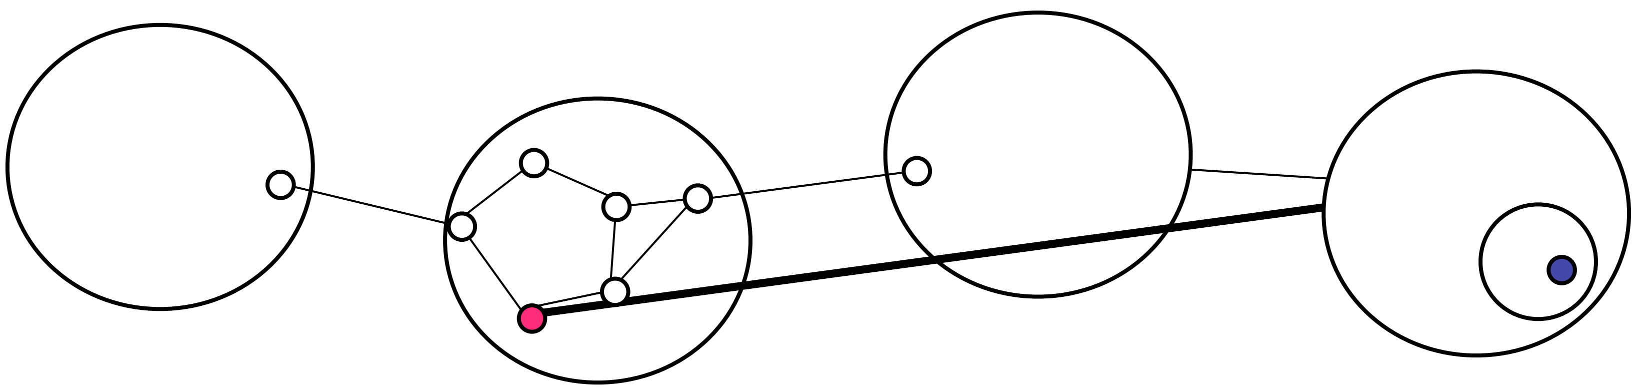

quasi-hierarchical

Forwarding Information

Quasi-hierarchical routing:

- next-hop node to reach any child of any ancestral aggregate
- if shortest path between two nodes in an aggregate remains within that aggregate, then finds shortest path from node to child aggregate of lowest-common ancestral aggregate containing source and destination

Strict-hierarchical routing:

- next-hop aggregate to reach any child of any ancestral aggregate, where next hop is sibling
- if shortest path between two children in an ancestral aggregate remains within that aggregate, then finds shortest path at level of children within ancestral aggregate
- look-ups per node bounded by (2 * levels) - 1, where levels is number of relevant ancestral aggregates of node

Types of Hierarchical Routing

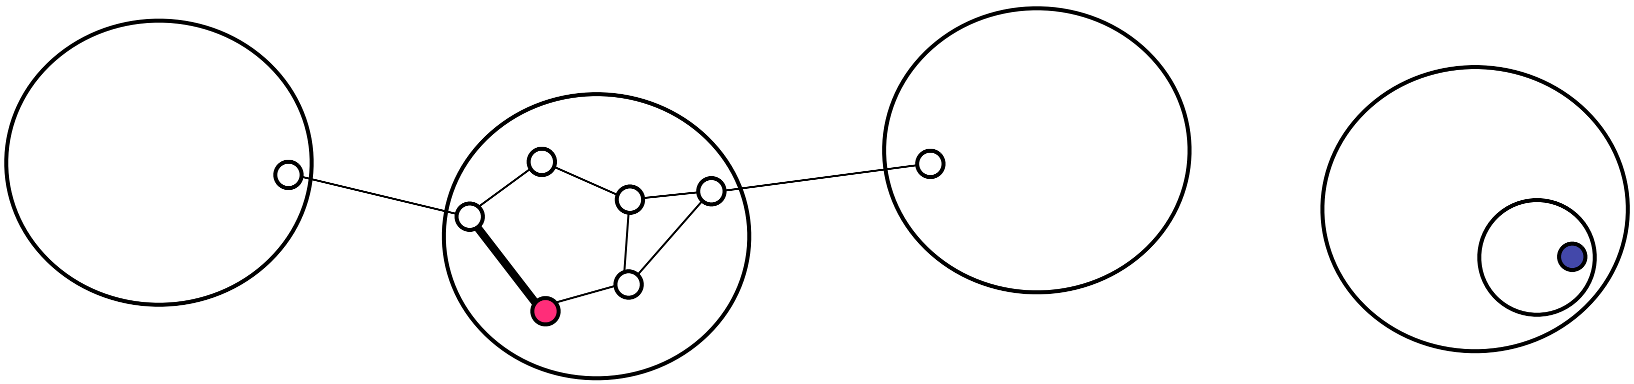

minimalist

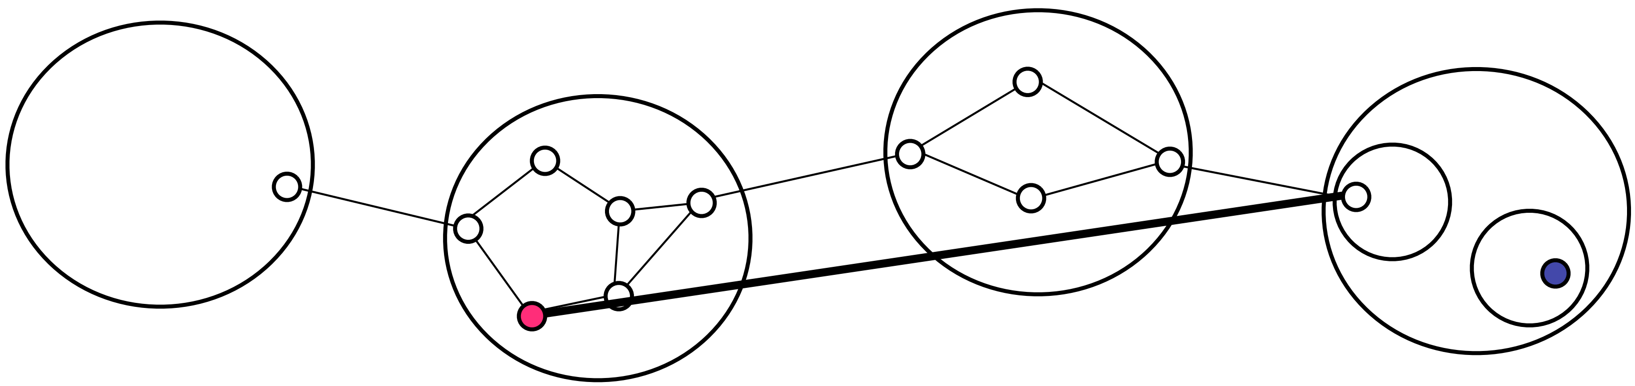

full monty

Forwarding Information

Minimalist hierarchical routing:

- for representative of ancestral aggregate, next-hop aggregate to reach any child of ancestral aggregate
- if shortest path between two children in an ancestral aggregate remains within that aggregate, then finds shortest path at level of children within ancestral aggregate
- simplest case: may route only to boundary of parent aggregate

Full monty:

- next-hop node to reach any node
- shortest path to any nodes
- may require many queries for sufficiently-detailed routing information

Nimrod as a Case Study

Routing is the primary motivation for the hierarchical network architecture

Designed for networks where ‘**who**’ and not ‘what’ is the focus of communication

Each aggregate has a name that is unique within its parent aggregate and an **address inherited from the aggregation hierarchy**

Each device has a **unique time-invariant name that is independent of its current address** derived from that of its associated lowest-level aggregate

Nimrod as a Case Study

Aggregation and abstraction algorithms as well as **route selection and packet forwarding algorithms** in use below the level of visible abstraction are chosen by the owners of the individual devices

Control functions (generation and distribution of state information for path generation, selection of paths and next hops, packet forwarding, and resolving names to addresses) **may be distributed among devices** according to the capabilities of devices and the responsiveness desired for the control functions

Accommodates unicast and **multicast** traffic as well as **mobile** devices

Nimrod as a Case Study

Route generation based on link states for flexible path choice permitting reconciliation of user needs and provider offerings and constraints

Link state at **multiple levels of abstraction** available to devices through automatic distribution and specific requests

Source-specified routing to accommodate private route selection criteria and successful packet forwarding in the presence of differing and potentially inconsistent views of state

Nested, concatenated virtual circuits for end-to-end packet forwarding along source-specified routes for efficiency and rapid local adaptation to changes in connectivity

Link State at Multiple Levels of Abstraction

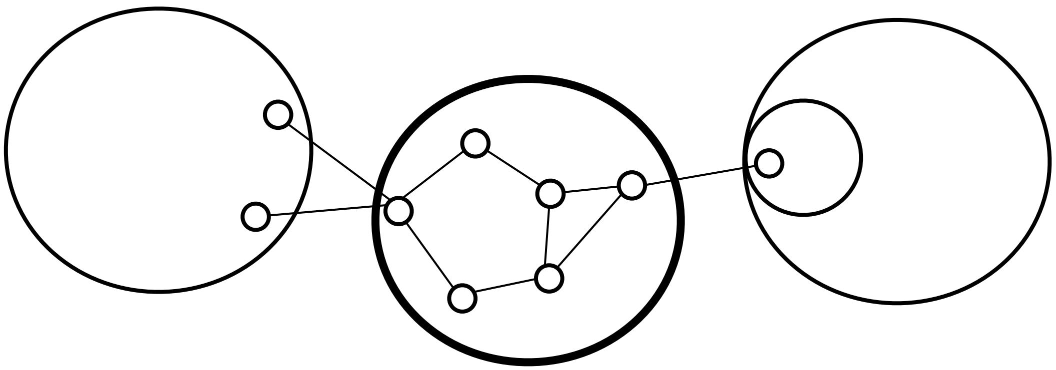

child components

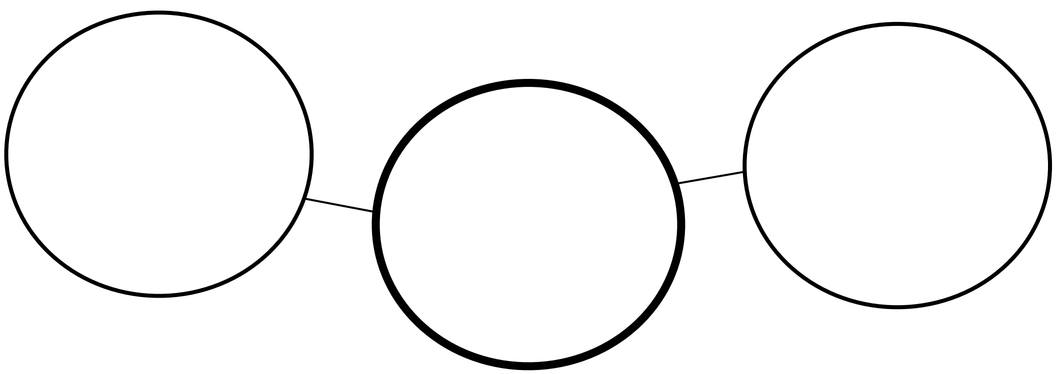

within parent aggregate

- what qualities of service
- between which entry and exit points
- at what times
- for which types of traffic
- between which users
- at what price

Examples of Topology Abstraction

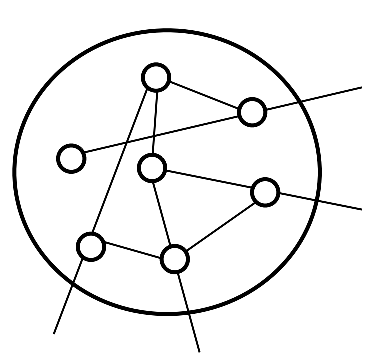

none

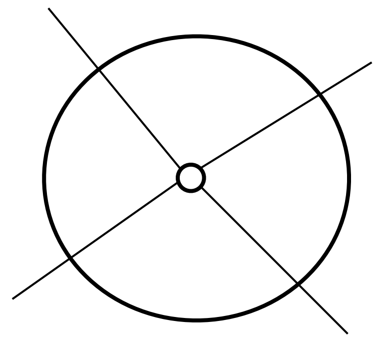

hub-spoke

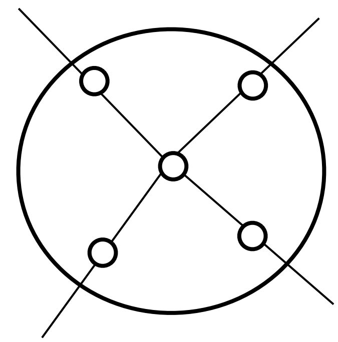

star

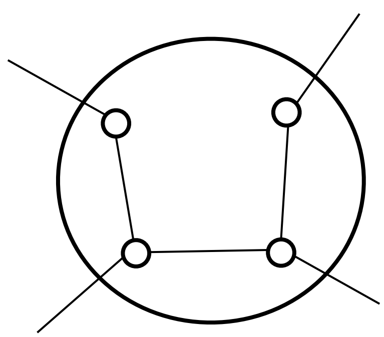

spanning tree

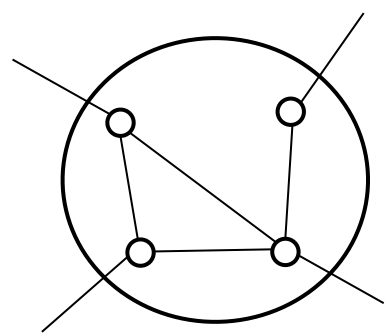

spanner

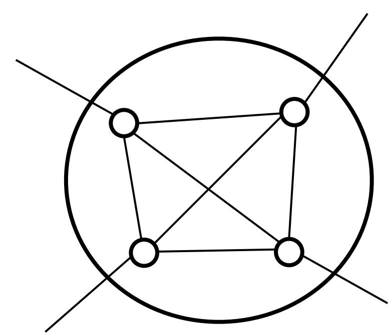

full mesh

Examples of Service Abstraction

Service restrictions:

- optimistic: union of service permissions
- pessimistic: intersection of service permissions

Services and fees:

- range of values
- average values
- mostly likely values given service restrictions

Flexibility of Virtual Circuits

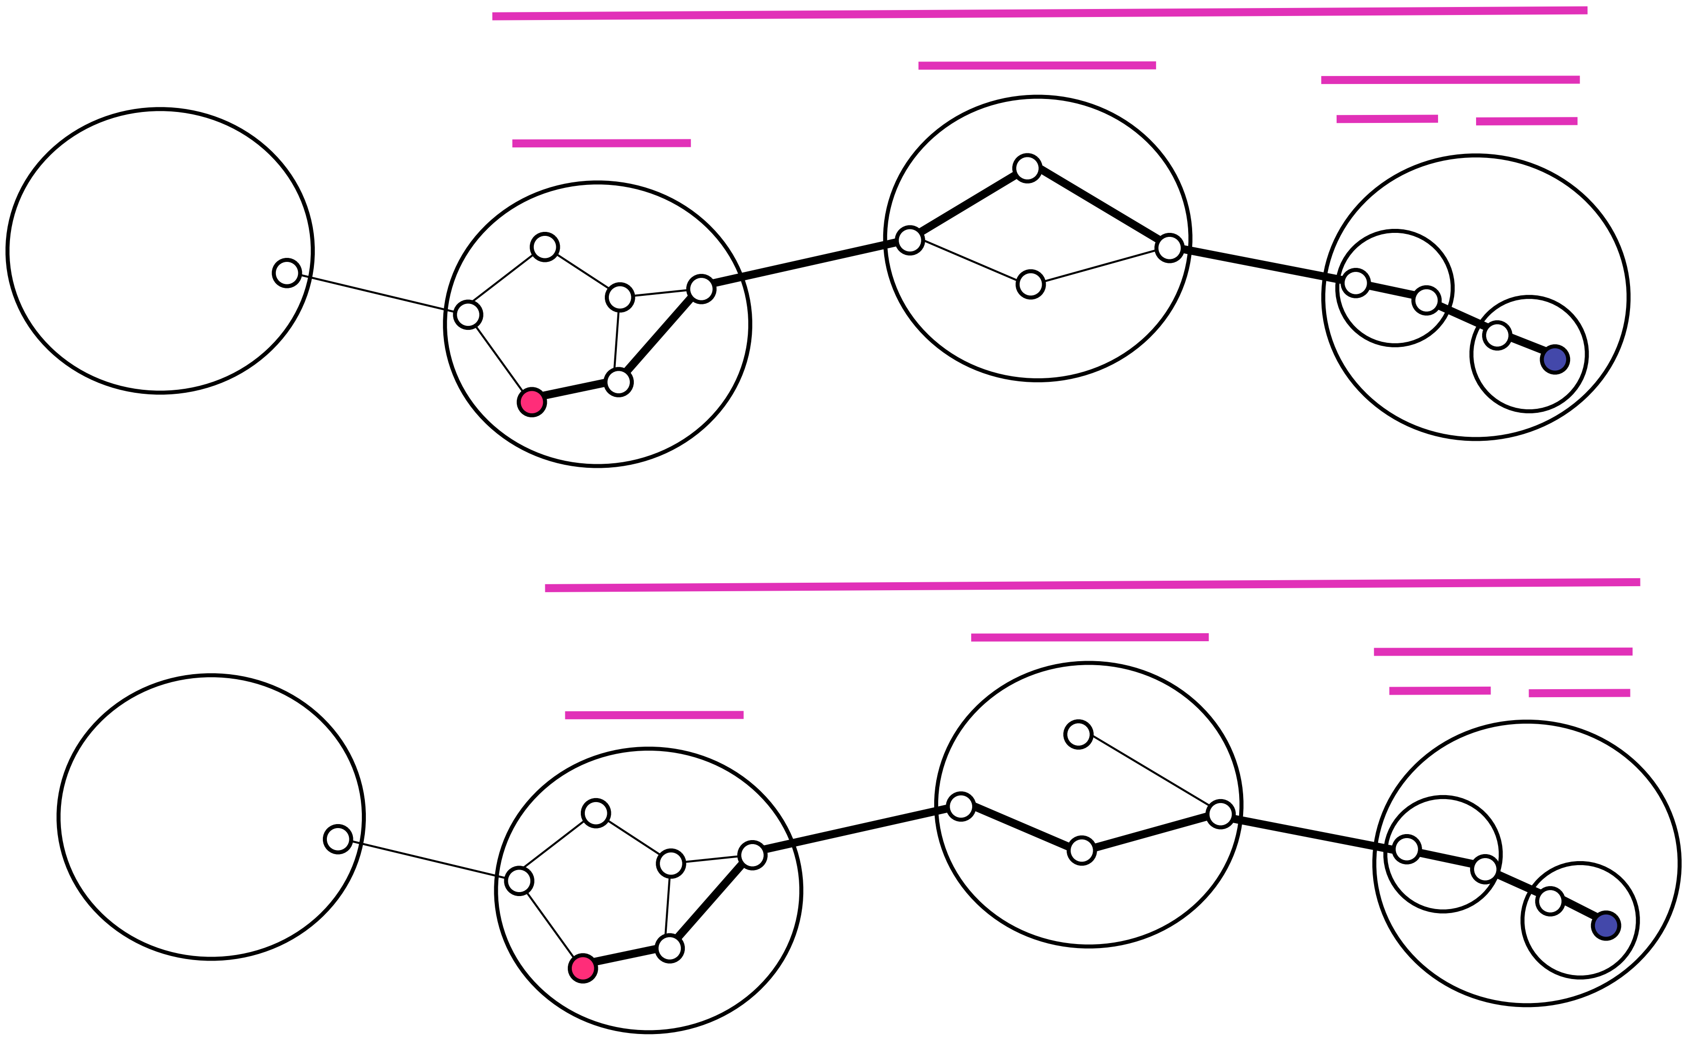

Important Trades to Consider

Degree of aggregation, granularity of abstraction, and specificity of source routes:

- probability of chosen path satisfying user and provider constraints
- transmission costs of distributing and requesting state and control information
- computation costs of aggregation, abstraction, and route selection

Important Trades to Consider

Automatic distribution of versus on-demand requests for state and control information:

- probability of using stale information to make control decisions
- delay-tolerance of routing functions
- transmission costs of distributing and requesting state and control information

Distribution of routing functions over devices:

- performance in terms of delay, throughput, and fault tolerance
- transmission costs of exchanging state and control information among devices
- computation costs for routing functions
- storage costs for state and control information

Applications as Drivers of Network Architecture

Characteristics:

- who is communicating and where
- what is being communicated
- number of senders and number of recipients
- quantity, delay-sensitivity, and loss-tolerance of information communicated

Alternatives to routing:

- flood search for information when caching is appropriate
- broadcast with network coding in environments with predominantly multipoint communications
- waiting for proximity of destinations before communication