

Survivable and Disruption Tolerant Networking

Issues, Challenges, and Research Directions

James P.G. Sterbenz

Lancaster University Computing Department, UK

University of Massachusetts Department of Computer Science, USA

jpgs@sterbenz.org

+1 508 944 3067

<http://www.sterbenz.org/jpgs/sumowin>

Rajesh Krishnan

BBN Technologies, USA

krash@bbn.com

Survivability & Disruption Tolerance

Abstract

This presentation surveys the issues and challenges in enhancing network survivability and providing application-to-application disruption tolerance, with emphasis on mobile wireless and long-delay communication. Conventional fault tolerance methods are necessary but not sufficient for survivability, since failures due to a coordinated attack are not random. Interlayer awareness (knobs and dials) is required for lower layers to convey their state upward, and for users and upper layers to exert influence downward.

A systematic approach is required on three levels:

1. Survivable network architectures that:
 - establish and maintain survivable topologies that strive to keep the network connected even under attack
 - design for communication in challenging environments in which the path from source to destination is not wholly available at any given instant in time; this requires new routing and forwarding mechanisms
 - the use of technology to enhance survivability such as adaptive networks and satellites
2. End-to-end (transport) protocols that are able to deal with weak, intermittent, and episodically connected and asymmetric channels, and are able to properly respond to channel-induced errors.
3. Adaptive applications that rely on knobs and dials from the network through the application to allow the user to influence the behaviour of the application based on network conditions, to best deal with and mask disruptions and delay.

We describe some of the issues and potential research directions to address each of these areas, and briefly present some example research projects.

Survivability & Disruption Tolerance

Outline

- Introduction to survivability and disruption tolerance
 - definitions, motivation, threats
 - environment: wireless, mobility, latency
 - challenges and assumptions
- Survivability and disruption tolerance strategy
 - network survivability
 1. maintain survivable connectivity when possible
 2. survivable communication even when not connected
 3. technologies to enhance survivability
 - end-to-end survivability and disruption tolerance
 - disruption-tolerant user-controlled adaptive applications
- Summary

Survivability & Disruption Tolerance

Definitions

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]

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]

Survivability & Disruption Tolerance

Motivation and Threats

- Network & applications should remain operational
 - when the network is under attack
 - particularly critical and lifeline services
note: 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

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
 - long path (c) or store-and-forward queueing
 - appears to be a disruption for interactive applications
 - 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

Survivability & Disruption Tolerance

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

Survivability & Disruption Tolerance

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

Survivability & Disruption Tolerance

Requirements and Goals

- Survivability goals for *network*
 - resistance to attack
 - recognition when attack has occurred
 - recovery from attack after occurrence
 - refinement in future response to attack
- 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

Survivability & Disruption Tolerance

Beyond Fault Tolerance and Cryptography

- Fault models do not hold under malicious attack
 - we cannot assume independence and random failures
 - therefore, fault tolerance is necessary, but *not* sufficient
- Cryptography does not ensure survivability
 - threat of traffic flow analysis leading to attack
 - when *data* encrypted *control* (signalling, headers) typically not
 - signal processing can extract information from encrypted data
 - control plane and physical infrastructure attacks
 - unauthenticated signalling
 - unencrypted packet headers

Survivability & Disruption Tolerance

Mobile, Wireless, and High Latency

- Traditional communication models
 - are not survivable even in wired network
 - redundancy not geographically diverse
 - adapt to episodic connectivity as faults ...
... that must be recovered
- Are not survivable...
 - when mobility exceeds reactivity of traditional control loops
 - when channel conditions don't allow end-to-end path
 - with links highly asymmetric or unidirectional
 - when routing protocols rarely or never converge

Survivability & Disruption Tolerance

Attacking the Problem

- New way of thinking needed
 - *expect* challenging channel environments
 - *expect and exploit* mobility
 - *expect, mask, and adapt to* high latency
- Attack problem at *all* levels
 - 1 robust physical channels
 - 2 robust links and survivable MAC
 - 3 survivable networks
 - 4 disruption tolerant end-to-end protocols
 - 7 adaptive user-controlled applications
- Need
 - interlayer awareness and control (knobs and dials)
 - intelligent resource & constraint tradeoffs (P, M, B, E, L)

Survivability & Disruption Tolerance

Network Survivability

- Introduction to survivability and disruption tolerance
- Survivability and disruption tolerance strategy
 - network survivability
 1. establish & maintain survivable connectivity when possible
 2. survivable communication even when not connected
 3. technologies to enhance survivability
 - end-to-end survivability and disruption tolerance
 - disruption-tolerant user-controlled adaptive applications
- Summary


Survivable Connectivity

Network Establishment

- Establishment of network structure and connectivity
 - secure auto-configuration and self-organisation
 - all infrastructure protocols and signalling must be
 - secure and resistant to attack
 - authenticated
 - use infrastructure when available ...
 - name servers
 - PKI, CA...
 - ... but don't depend on it: take local actions when necessary
- Maintain connectivity when practical
 - without sacrificing other requirements

Network Self-Organisation


Neighbour Discovery



- Nodes emit beacons to announce their presence
 - known frequencies and codes used for announcements
- Establishes set of directly reachable nodes

Network Self-Organisation

Link Formation

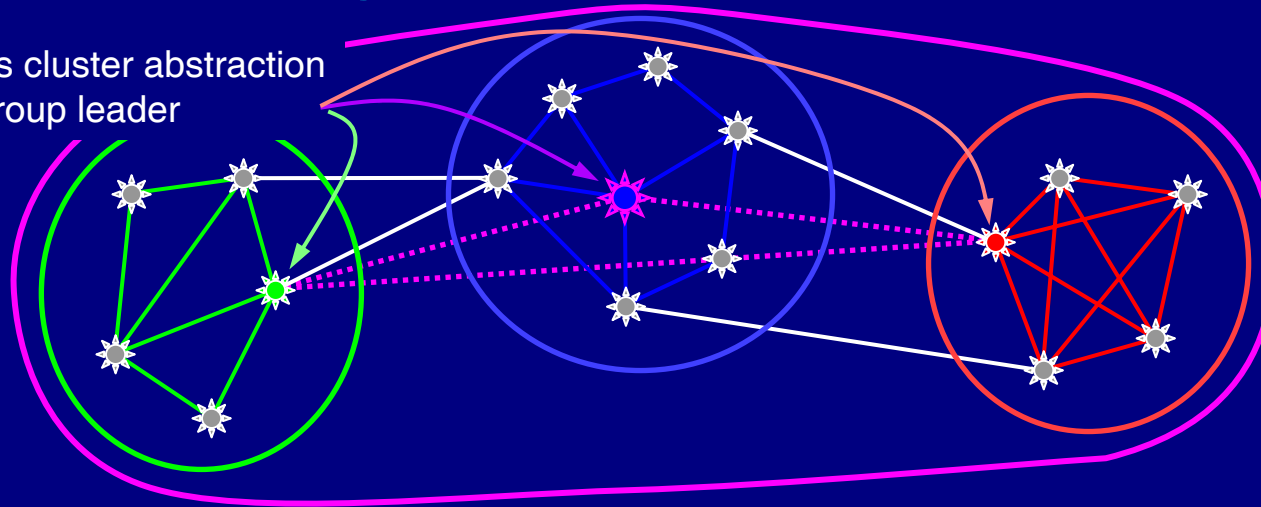


- Pairwise negotiation of link formation
 - interested nodes answer beacons
 - exchange identification, node and link characteristics
 - layer 2 connectivity structure
- Maintain link adjacencies
 - e.g. **keepalive** messages

Network Self-Organisation

Self-Organisation and Federation


leaderless cluster abstraction
or peer group leader



- Communicating nodes self-organise into federations
 - address acquisition
 - hierarchical cluster formation and leader election
 - based on administrative concerns, security, role/task based
 - bootstrap routing topology

Network Self-Organisation

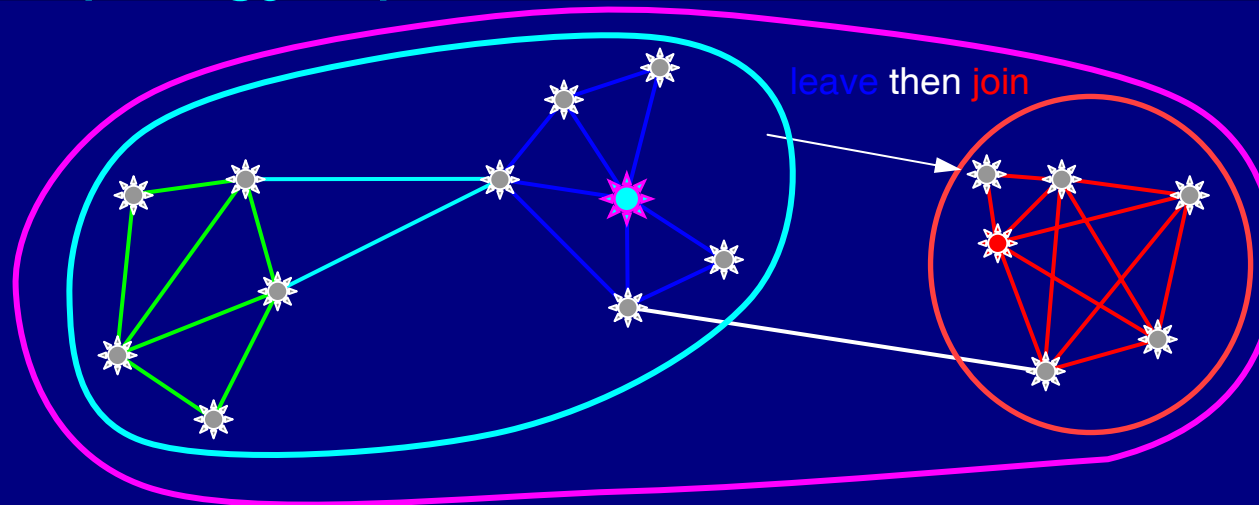
Topology Optimisation and Maintenance



- Topology maintenance of federations
 - merge/split
 - group mobility, dynamic coalitions
 - heal partition

Network Self-Organisation

Topology Optimisation and Maintenance



- Topology maintenance of nodes
 - node mobility
 - leave/join from/to federation
 - resolution to identifier vs. topological address reassignment

Survivable Connectivity


Establishment and LPD

- Low probability of detection (LPD)
 - low transmission power to limit detection
 - stealthy network is more resistant to attack ...
... but stealth makes legitimate communication difficult

Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission Power



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity



Survivable Connectivity

Topological Connectivity: Transmission Power

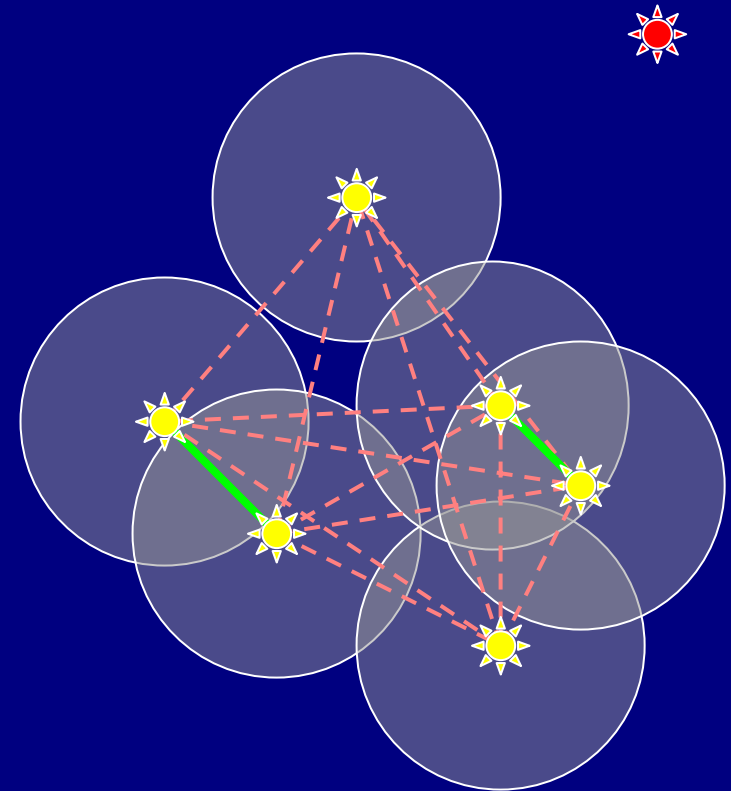
- Transmission power
 - low:
 - no connectivity



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands



Survivable Connectivity

Topological Connectivity: Transmission Power

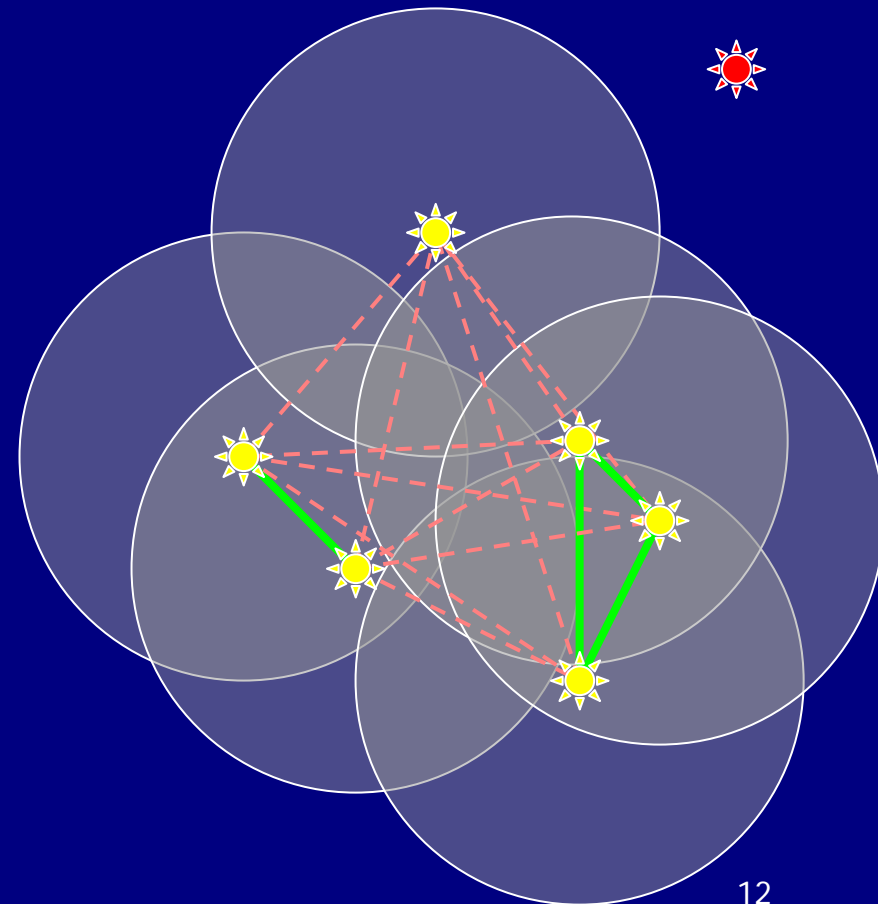
- Transmission power
 - low:
 - no connectivity
 - partitioned islands



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands



Survivable Connectivity

Topological Connectivity: Transmission Power

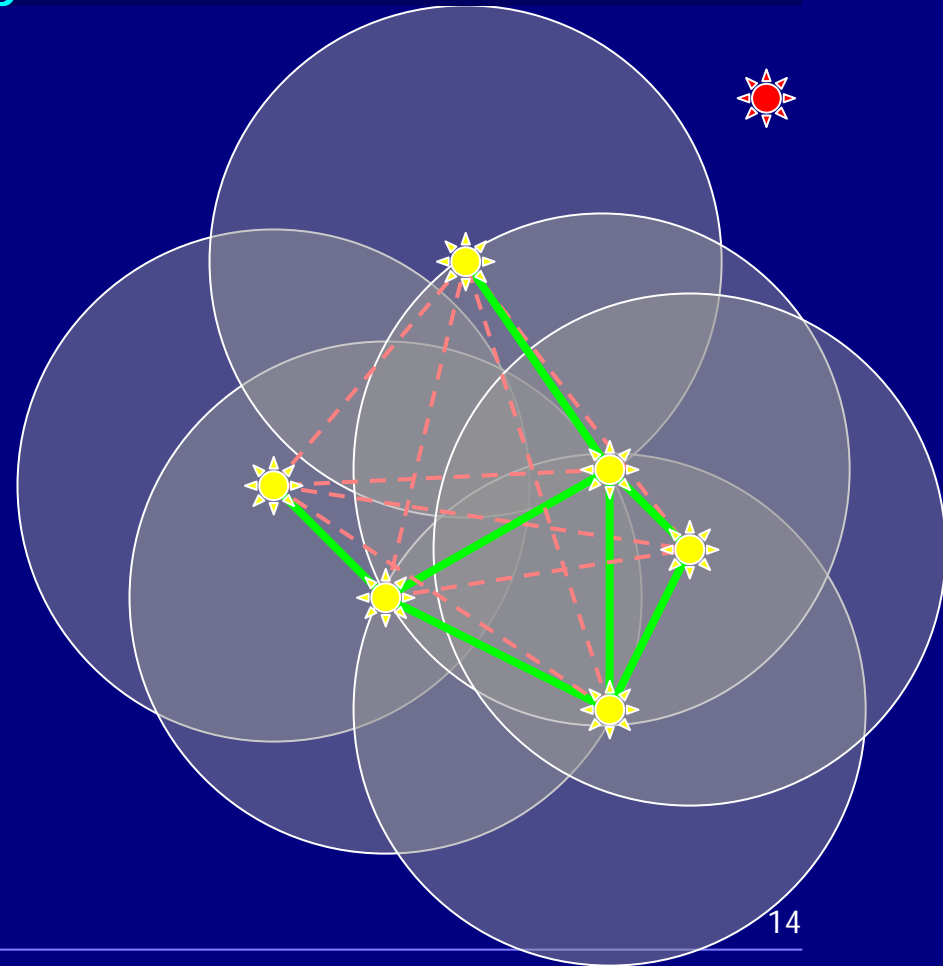
- Transmission power
 - low:
 - no connectivity
 - partitioned islands



Survivable Connectivity

Topological Connectivity: Transmission Power

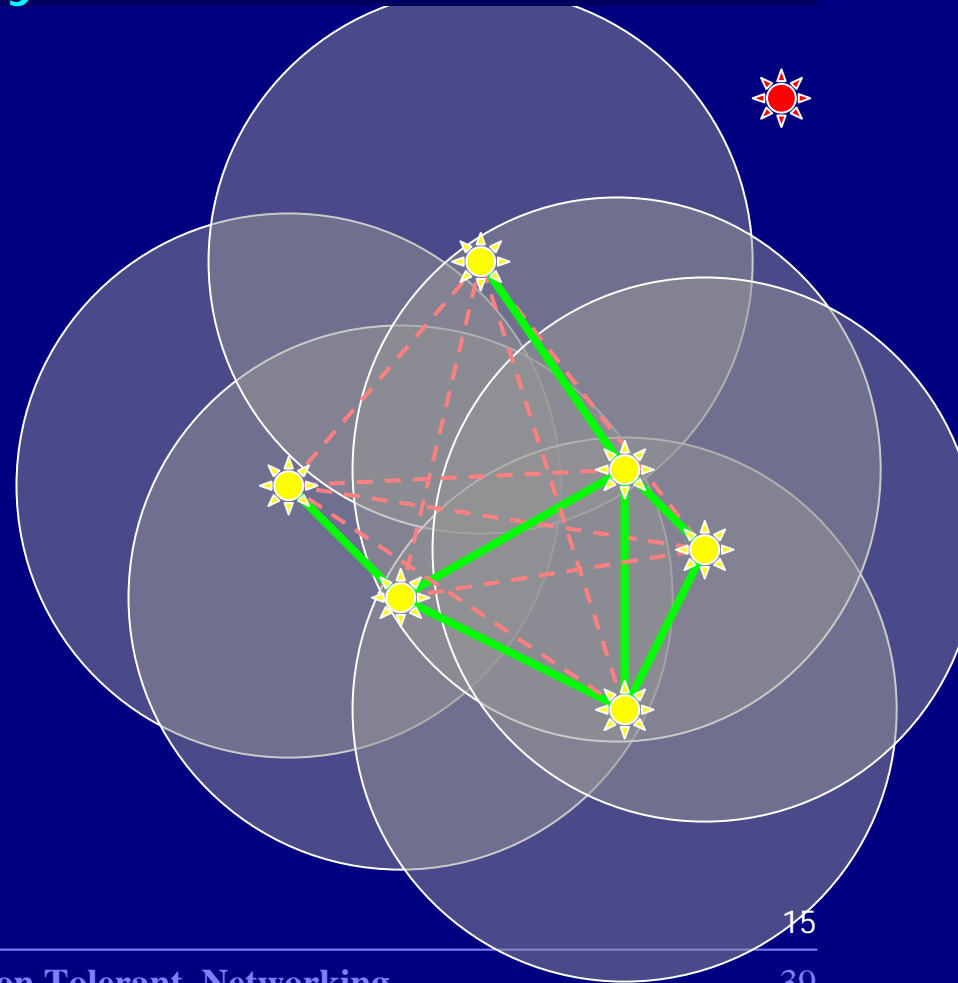
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected



Survivable Connectivity

Topological Connectivity: Transmission Power

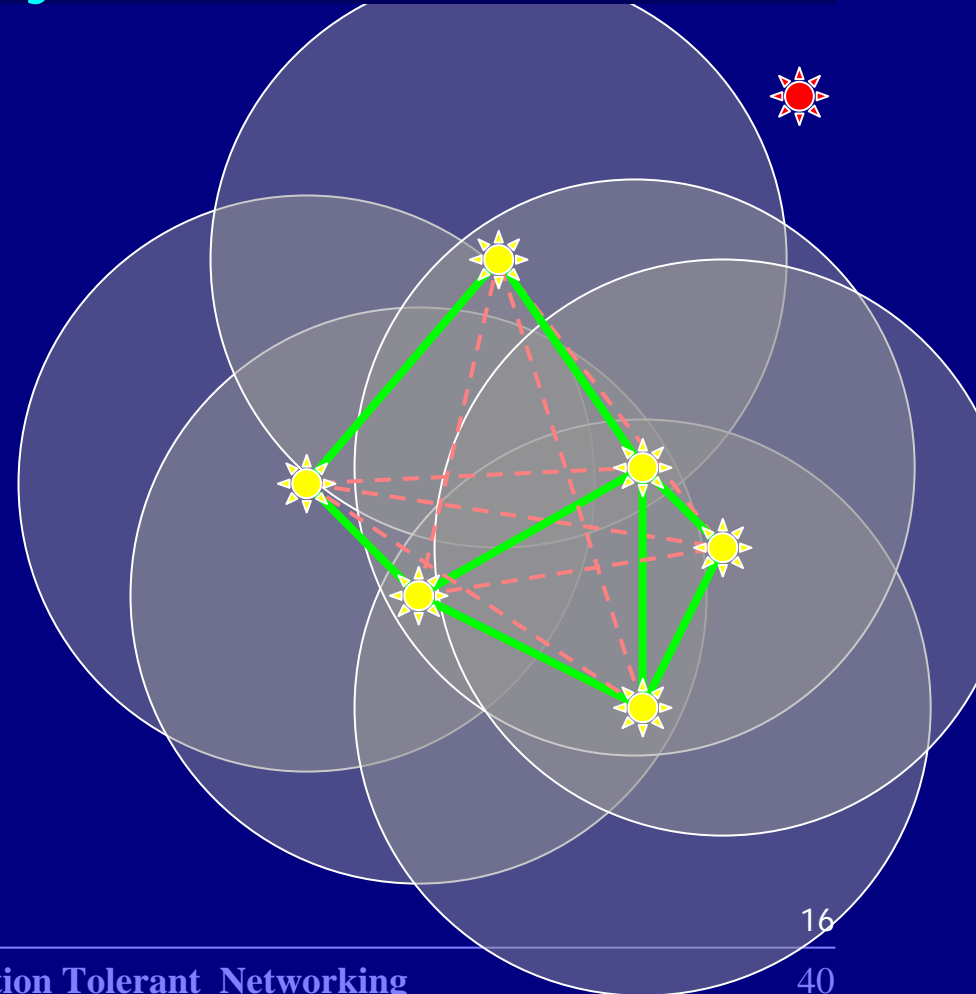
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected



Survivable Connectivity

Topological Connectivity: Transmission Power

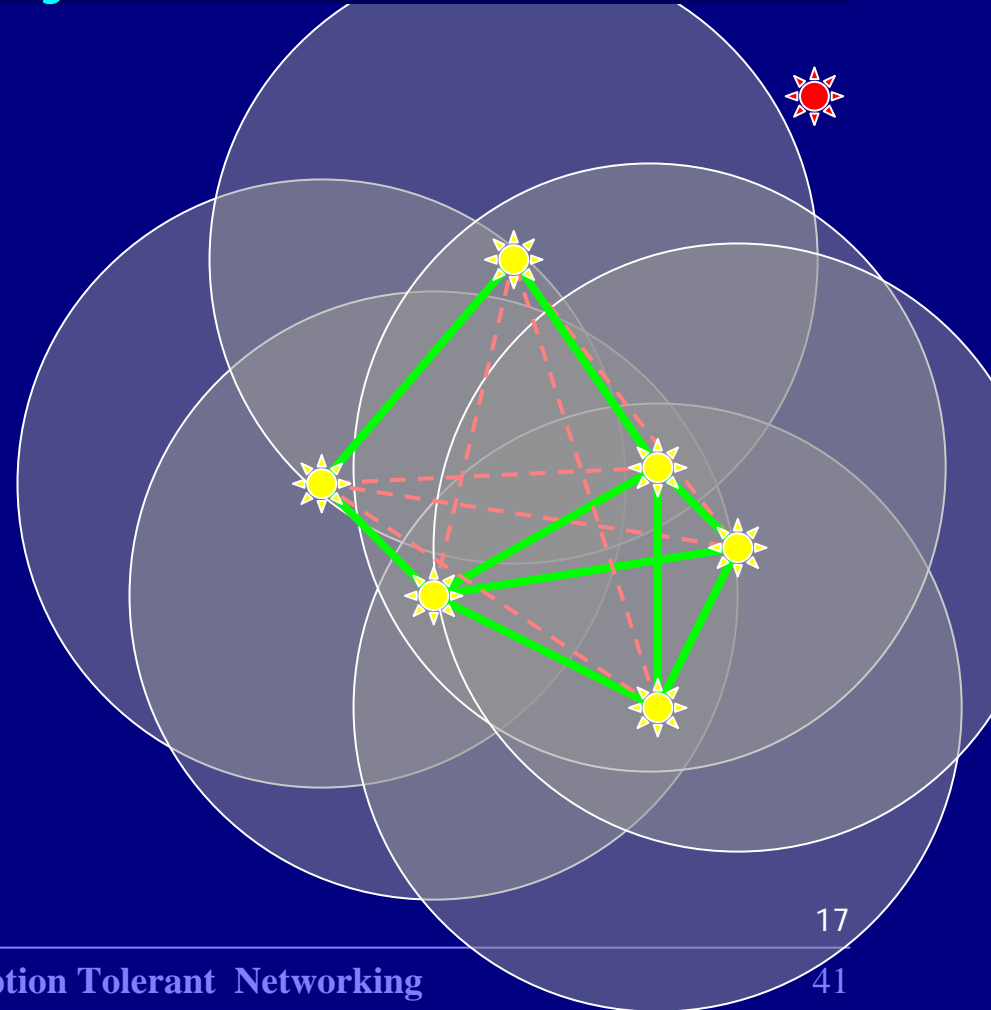
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected



Survivable Connectivity

Topological Connectivity: Transmission Power

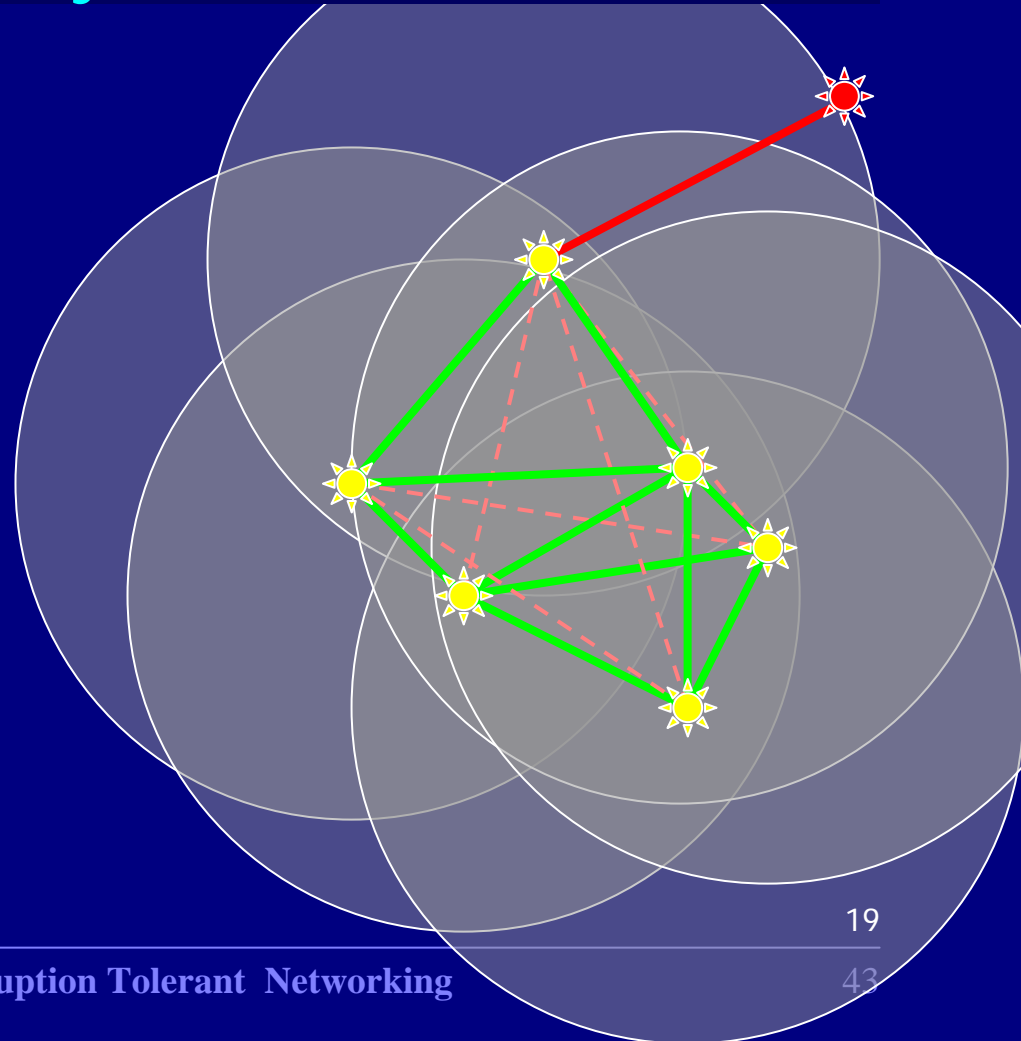
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power

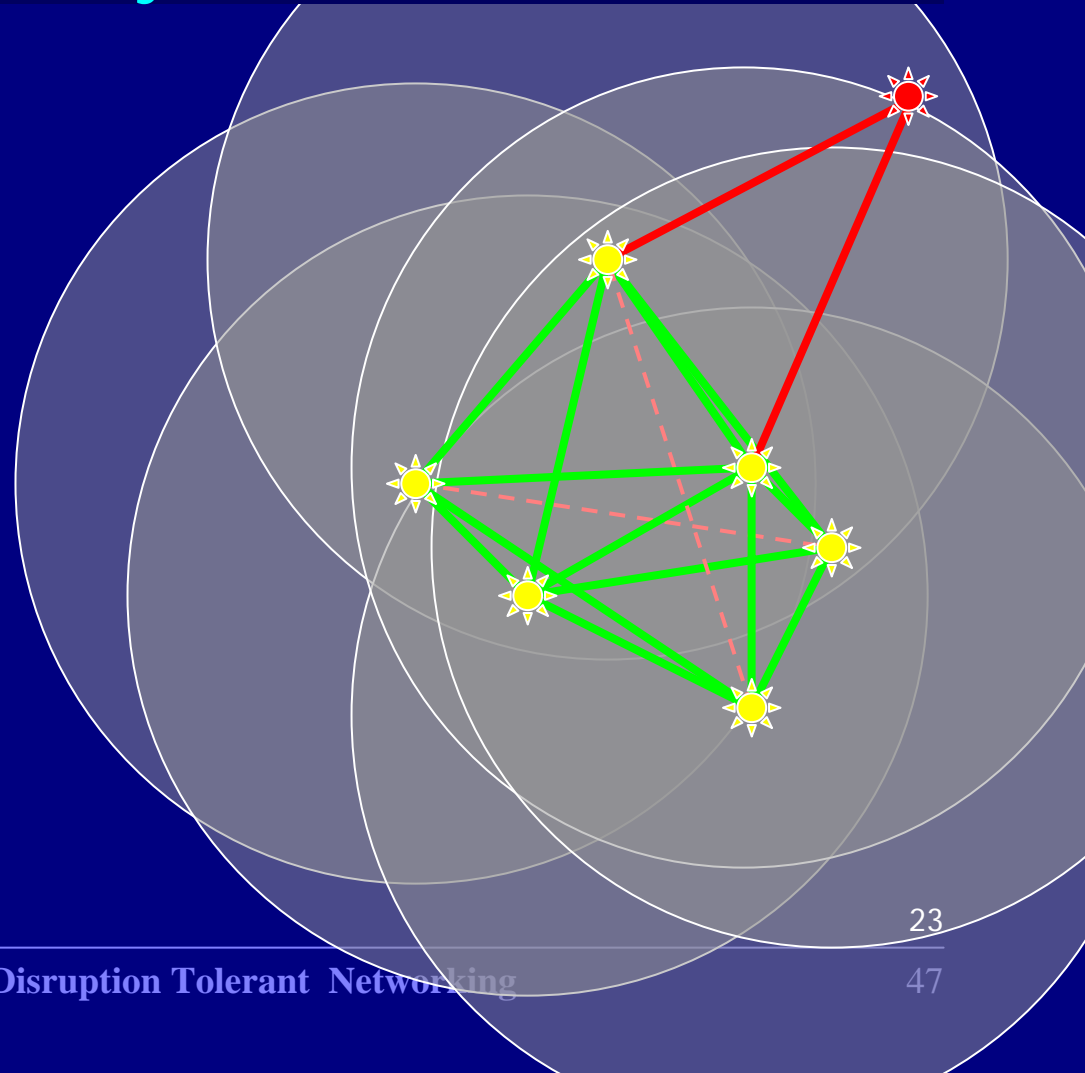
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power

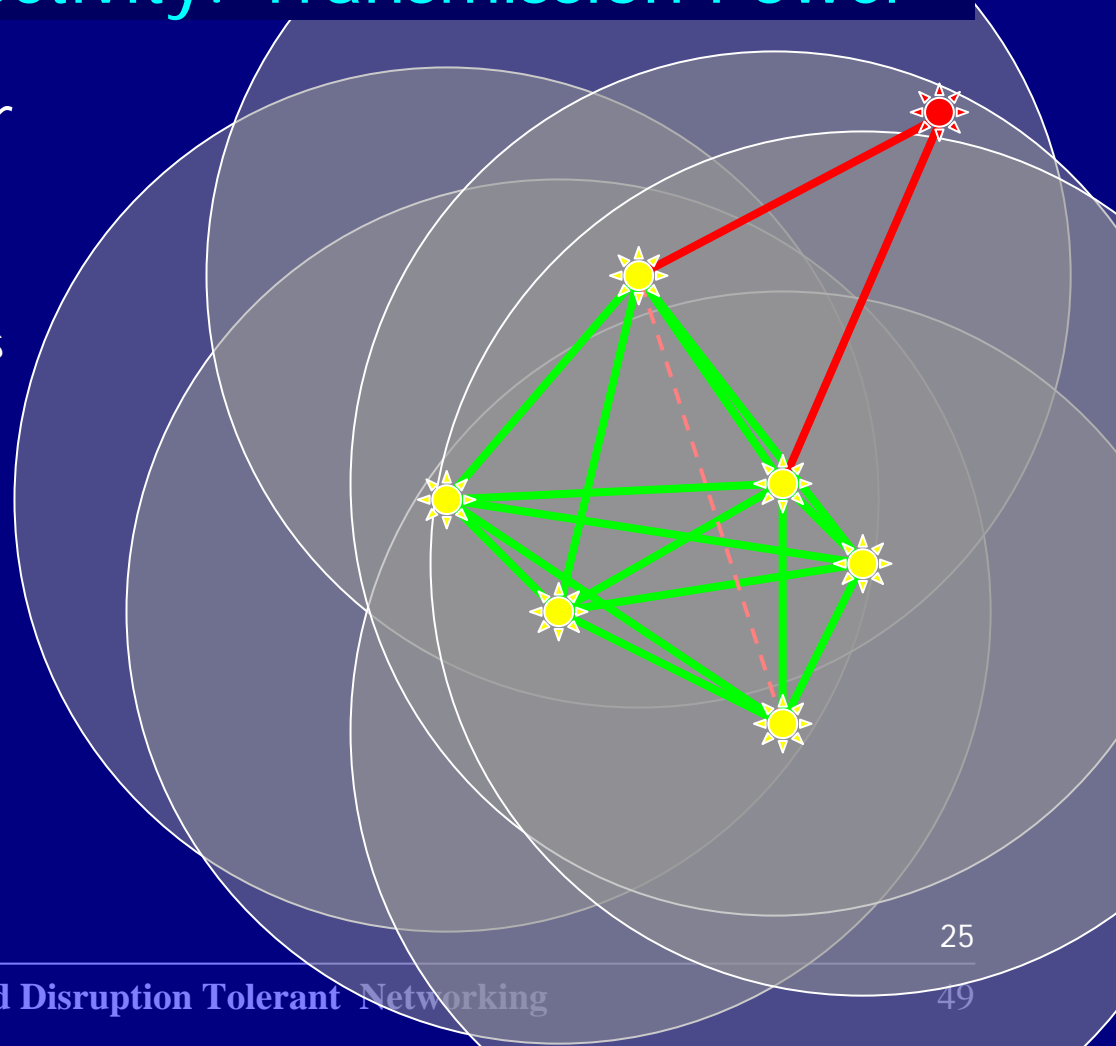
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power


- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power

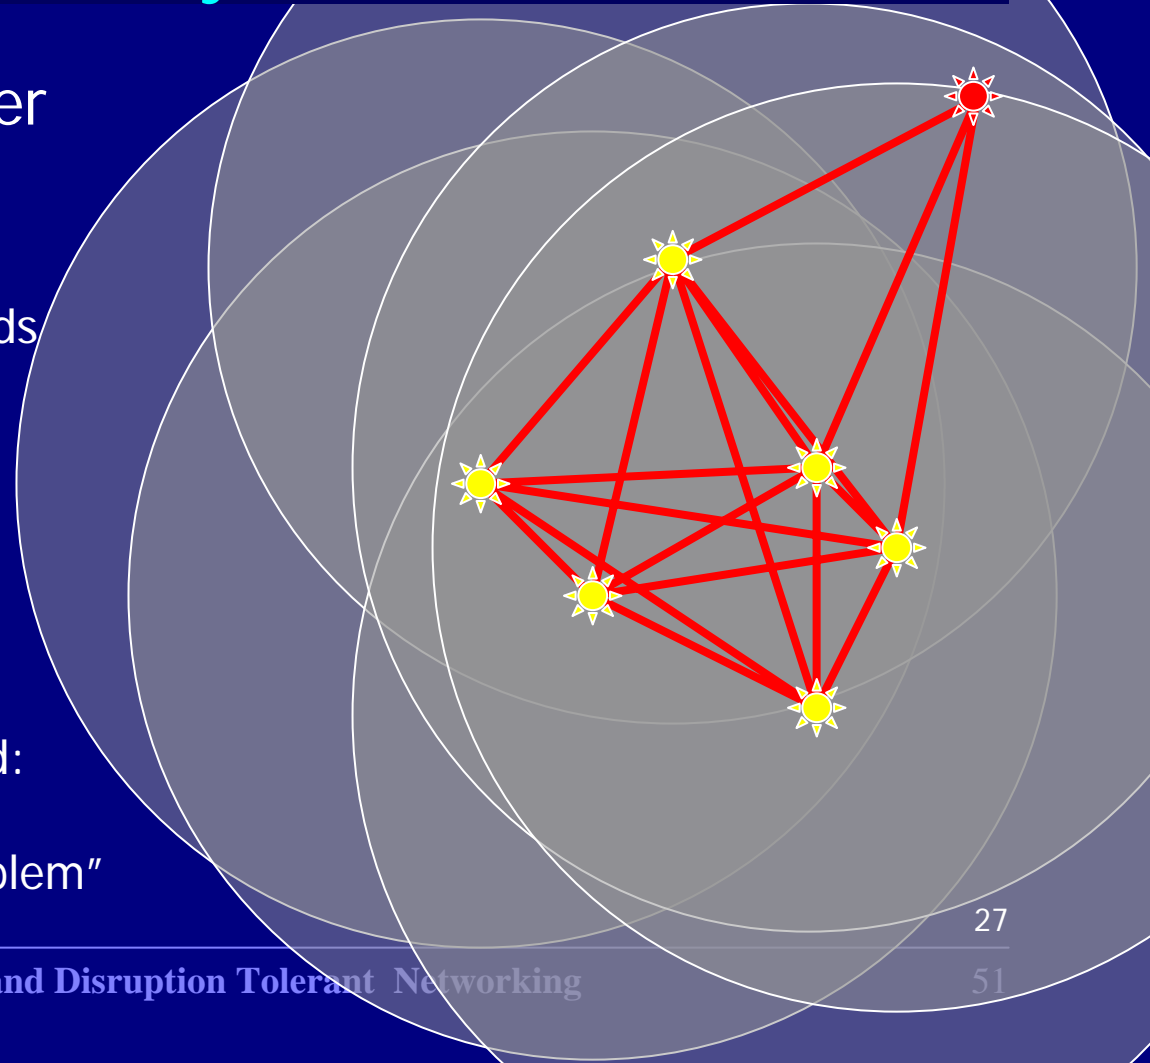
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth



Survivable Connectivity

Topological Connectivity: Transmission Power

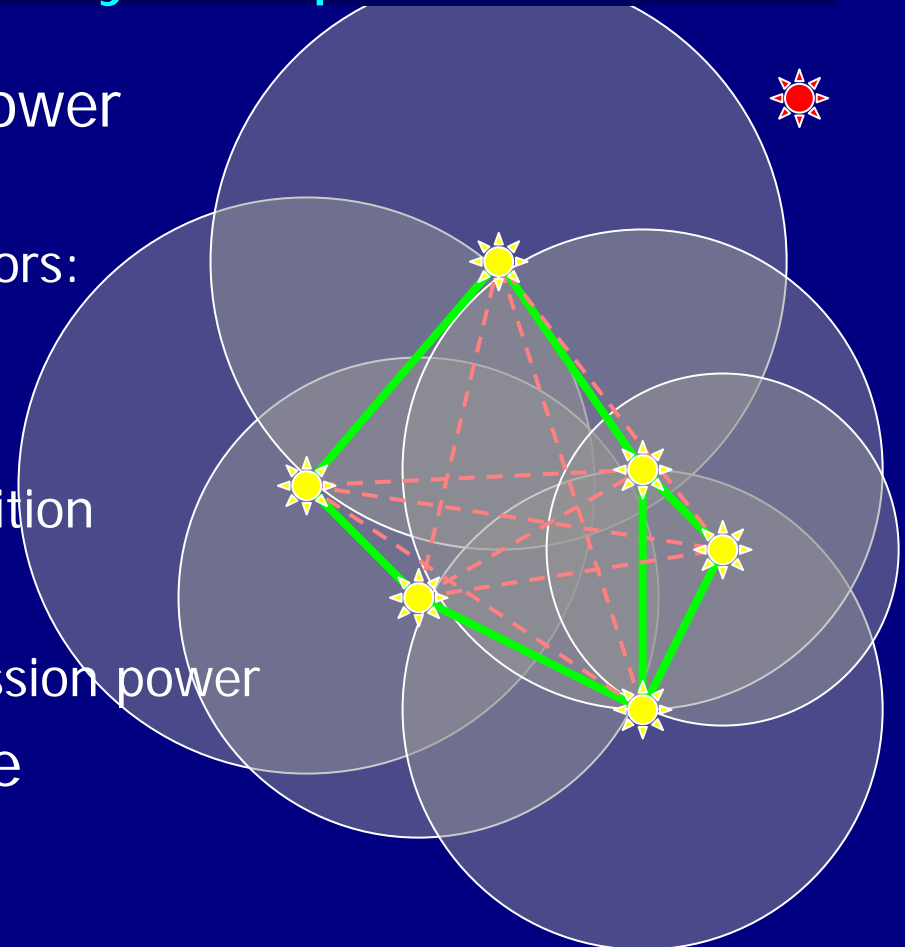
- Transmission power
 - low:
 - no connectivity
 - partitioned islands
 - sufficient
 - connected
 - biconnected
 - excessive
 - lack of stealth
 - highly connected:
self jamming
“parking lot problem”



Survivable Connectivity

Topological Connectivity: Adaptive Power

- Adaptive transmission power
 - each node adjusts
 - control number of neighbors:
degree of connectivity
- Biconnected graph
 - single link cut avoids partition
- May be more stealthy
 - in cases of lower transmission power
- Omnidirectional antennæ

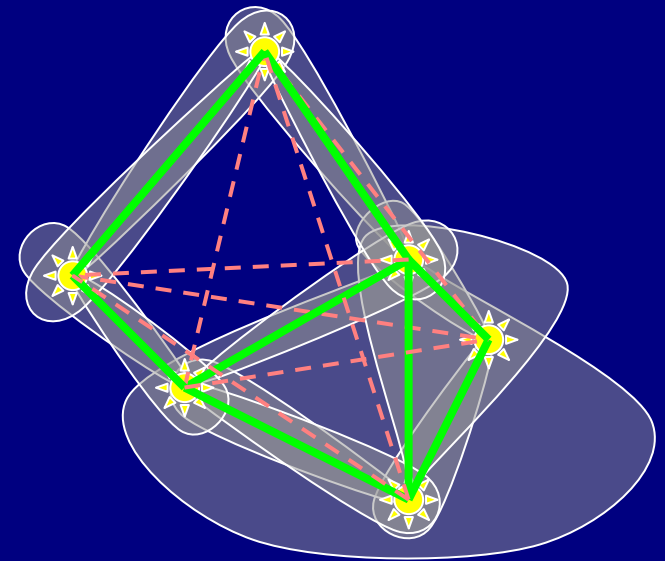


Survivable Connectivity

Topological Connectivity: Directional Antennæ




- Directional antennæ focus
 - transmission into sector
 - increase spatial reuse
- Reduced transmission
 - with better connectivity
- Increased complexity in:
 - antenna design
 - node discovery
 - MAC protocols (steering)
 - mobility tracking



Survivable Connectivity

Topological Connectivity: Directional Antennæ

- Directional antennæ
 - focus transmission into sector
 - increase spatial reuse
- Reduced transmission power
 - with better connectivity
- Increased complexity in:
 - antenna design
 - node discovery
 - MAC protocols (steering)
 - mobility tracking
- Increased stealth
 - assuming receiver locations known



Network Survivability Strategy

Survivable Communication

- Introduction to survivability
- Survivability strategy
 1. establish/maintain survivable connectivity when possible
 2. survivable communication even when not well connected
 - expect weak and episodic connectivity
 - expect and exploit mobility
 3. technologies to enhance survivability
- Summary

Communication Background

Network Layer Service and Interfaces

- Network layer 3 is above link layer 2
 - *addressing* : network layer identifier for end systems (hosts)
 - *forwarding* : transfers packets hop-by-hop
 - using link layer services
 - network layer responsible for determining *which* next hop
 - *routing* : determination of path to forward packets
 - *signalling* : messages to control network layer behaviour
 - *traffic management* : management of traffic and congestion
- Network layer service to transport layer (L4)
 - deliver TPDU's to destination transport entity

Communication Background

Forwarding vs. Routing

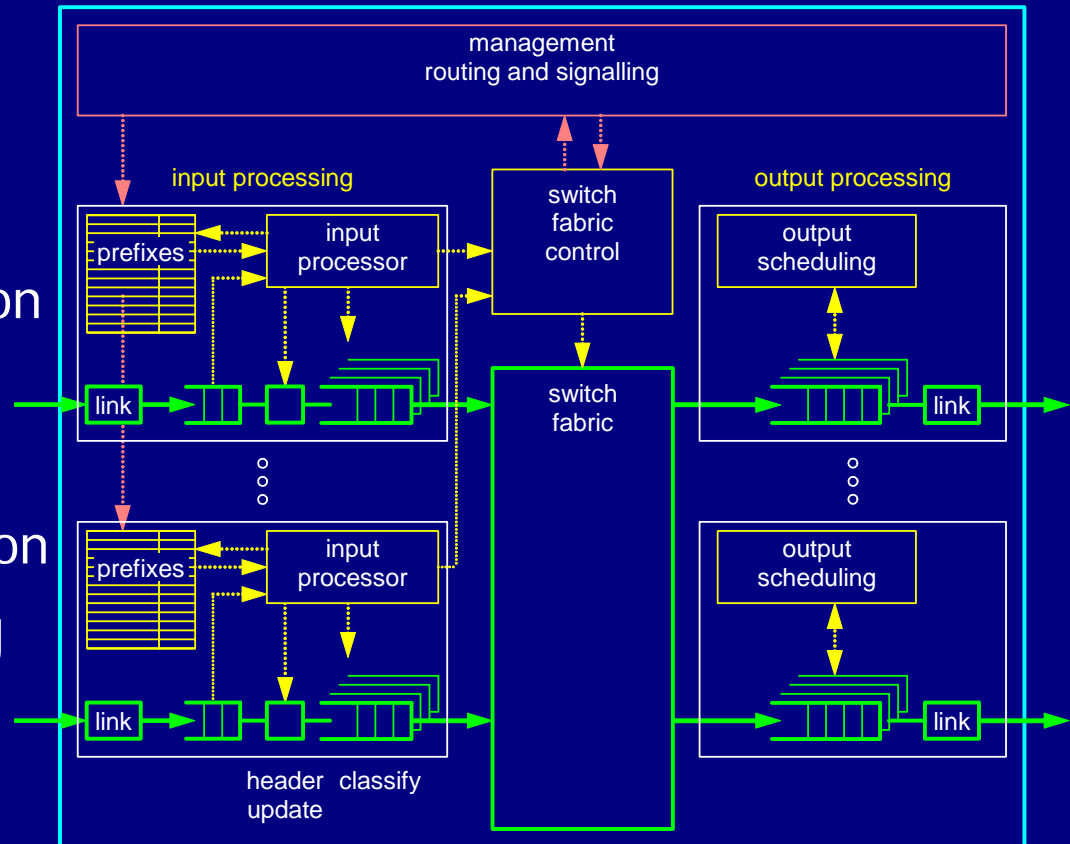
- *Forwarding* transfers packets hop-by-hop
 - each switch (router) makes decision on which link to send
 - forwarding table (generally) used to make decision
 - forwarding is *per packet* decision
 - [analogy: determining which exits to take on a drive]
- *Routing* determines the path to take
 - routing algorithm [lecture R] independent of forwarding
 - forwarding table entries populated by routing
 - routing is (generally) not done per packet
 - [analogy: planning trip from source to destination]

Forwarding and routing are very different

Communication Background

Fast IP Datagram Switch Architecture

- Routing
 - loads prefix table
- Input processing
 - packet classification
 - IP lookup
- Switch fabric
 - port interconnection
- Output processing
 - packet scheduling




Non-Survivable Communication

Routing Convergence and Mobility

- Current routing algorithms assume *eventual stability*
 - converge to stable communication paths
 - complete end-to-end path must exist at some point in time
 - link outage treated as *fault* that must be repaired
- Moderate mobility is *tolerated* as a topology change

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- Among possible links

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- Among possible links
 - network is formed
 - biconnected if possible

Non-Survivable Communication

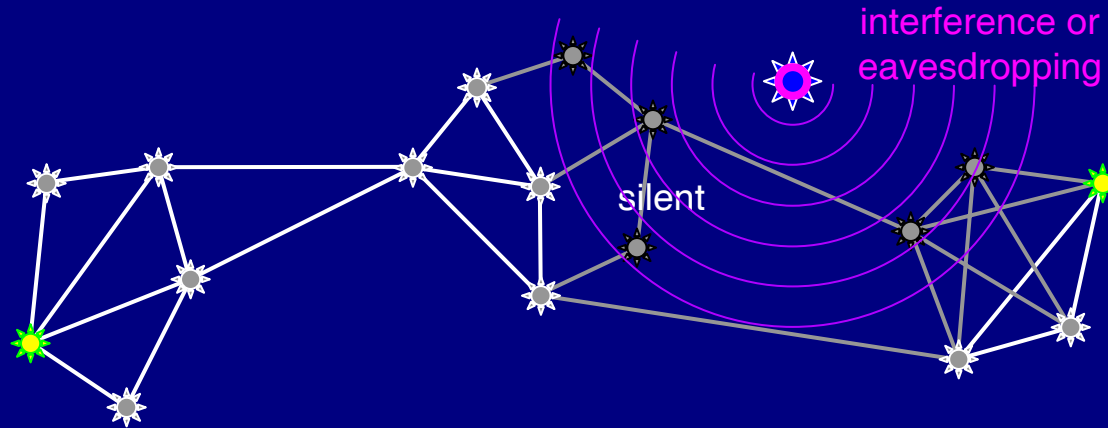
Eventual Stability: Wait for Complete Path



- Among possible links
 - network is formed
 - biconnected if possible

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While **interference** or **suspected eavesdropping**...
 - routing can't converge on a source → destination path

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source → destination path

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source \rightarrow destination path
- Routing algorithms recompute and converge

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source → destination path
- Routing algorithms recompute and converge
 - (complete) source → destination path exists

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source → destination path
- Routing algorithms recompute and converge
 - (complete) source → destination path exists
 - **data can be transferred along path** (as long as stable)

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source → destination path
- Routing algorithms recompute and converge
 - (complete) source → destination path exists
 - **data can be transferred along path** (as long as stable)

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source → destination path
- Routing algorithms recompute and converge
 - (complete) source → destination path exists
 - **data can be transferred along path** (as long as stable)

Non-Survivable Communication


Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source → destination path
- Routing algorithms recompute and converge
 - (complete) source → destination path exists
 - **data can be transferred along path** (as long as stable)

Non-Survivable Communication

Eventual Stability: Wait for Complete Path



- While interference or suspected eavesdropping...
 - routing can't converge on a source → destination path
- Routing algorithms recompute and converge
 - (complete) source → destination path exists
 - **data can be transferred along path** (as long as stable)


Survivable Communication

Routing Convergence

- Need to *assume* weak and episodic connectivity
 - routine occurrence for which network is designed
- Survivable communication: *eventual connectivity*
 - communicate as far as possible, whenever possible
 - hold data when necessary (store-and-forward)
 - deflection when necessary (buffer limitations)
 - schedule transmission for optimum LPI/LPD and energy
 - optimise for eventual stability when possible
 - avoid store-and forward avoidance
 - when stable path *is* available
 - cut-through switches

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- **Multiple interferences or suspected eavesdroppers**
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - **store-and-forward when necessary**

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication


Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary

Survivable Communication

Eventual Connectivity



- Multiple interferences or suspected eavesdroppers
 - prevent an end-to-end path from *ever* existing
 - transfer data as far as possible
 - store-and-forward when necessary


Survivable Communication

Expect Mobility

- Routing and forwarding *expect* mobility
 - use location and trajectory information when available
 - direct information to expected location

Survivable Communication


Expect Mobility



- Routing and forwarding expect mobility
- Use location/trajectory information where available
 - unicast when predictable (e.g. planetary or racetrack UAV)

Survivable Communication


Expect Mobility



- Routing and forwarding expect mobility
- Use location/trajectory information where available
 - unicast when predictable (e.g. planetary or racetrack UAV)

Survivable Communication


Expect Mobility



- Routing and forwarding expect mobility
- Use location/trajectory information where available
 - unicast when predictable (e.g. planetary or racetrack UAV)
 - multicast to **area of expected location** (spray routing)

Survivable Communication

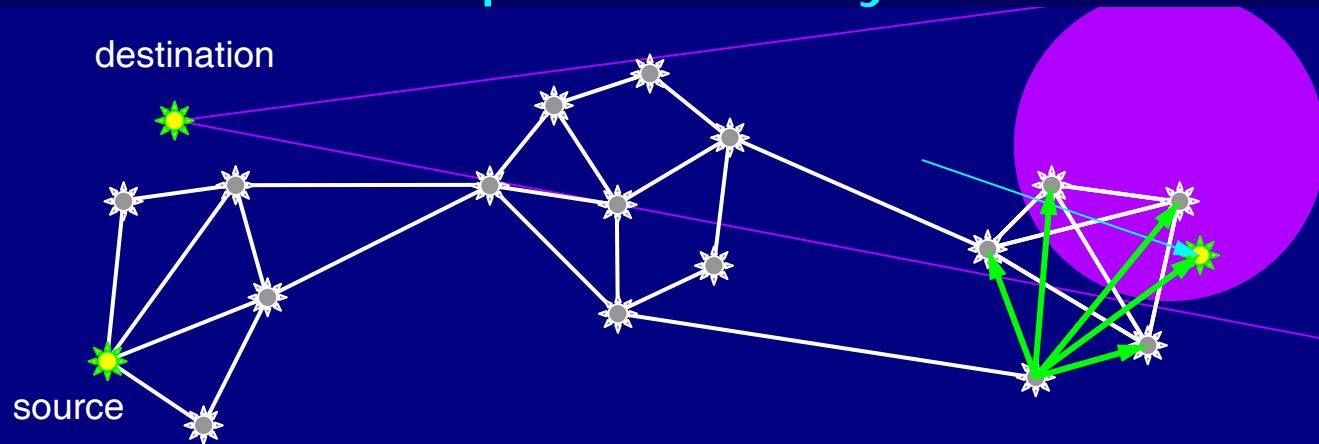
Expect Mobility



- Routing and forwarding expect mobility
- Use location/trajectory information where available
 - unicast when predictable (e.g. planetary or racetrack UAV)
 - multicast to **area of expected location** (spray routing)

Survivable Communication

Expect Mobility



- Routing and forwarding expect mobility
- Use location/trajectory information where available
 - unicast when predictable (e.g. planetary or racetrack UAV)
 - multicast to **area of expected location** (spray routing)
 - cluster may have **inherent broadcast or epidemic routing**


Survivable Communication

Exploit Mobility

- Position node/antenna for survivability
 - use trajectory information when available
 - exert control on movement of other nodes
- Node can carry data as they move
 - *store-and-haul* data without radiating transmissions
 - transit areas of *no* channel connectivity

Survivable Communication


Exploit Mobility



- Multiple interferences or suspected eavesdroppers

Survivable Communication


Exploit Mobility



- Multiple interferences or suspected eavesdroppers

Survivable Communication


Exploit Mobility



- Multiple interferences or suspected eavesdroppers

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication

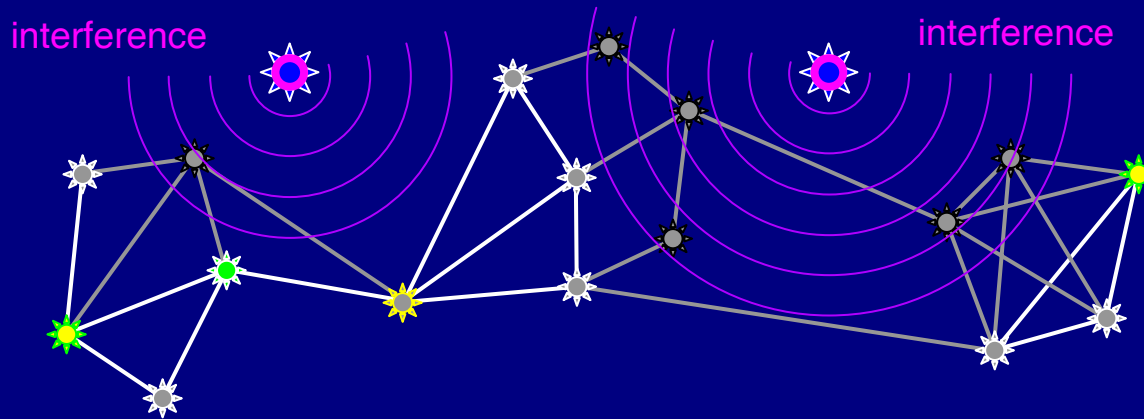
Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference

Survivable Communication

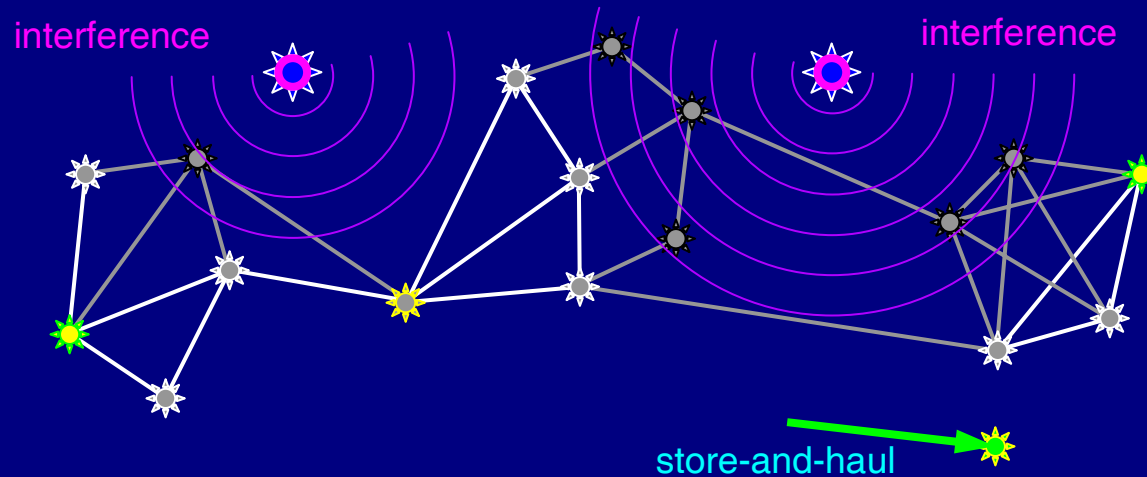
Exploit Mobility



- Move nodes and steer antenna around interference
- Mobile nodes haul data without radiating
 - interference and adversary node avoidance
 - transit disconnectivity

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference
- Mobile nodes haul data without radiating
 - interference and adversary node avoidance
 - transit disconnectivity

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference
- Mobile nodes haul data without radiating
 - interference and adversary node avoidance
 - transit disconnectivity

Survivable Communication


Exploit Mobility



- Move nodes and steer antenna around interference
- Mobile nodes haul data without radiating
 - interference and adversary node avoidance
 - transit disconnectivity

Survivable Communication

Exploit Mobility



- Move nodes and steer antenna around interference
- Mobile nodes haul data without radiating
 - interference and adversary node avoidance
 - transit disconnectivity

Survivable Communication

Adjust Data Transfer to Knowledge

- Opportunistic ...
 - epidemic routing protocols
 - transfer data when links are available and nodes reachable... but scoped and scheduled to:
 - reduce load while maintaining probability of delivery
 - reduce offered load to network while maintaining goodput
- Exert control on:
 - node and subnetwork movement
 - protocol and parameter choices
 - layer 2 connectivity and layer 3 federation topology

Opportunistic worst case bound; exploit knowledge to improve

Survivable Communication

Adjust Data Transfer to Environment

- Cut-through (when stable path available)
 - lowest latency for nodes that are capable
 - exploit “traditional” physical layer techniques
- Store-and-forward
 - immediate when link available to next node & empty queues
 - move data burst to other nodes for load balancing
- Store and forward with scheduled transfer
 - wait until link available to next node
 - new physical layer opportunities for burst transfer
- Store-and-haul data

Design for eventual connectivity, optimize for eventual stability

Network Survivability Strategy

Survivability Technologies

- Introduction to survivability
- Survivability strategy
 1. establish/maintain survivable connectivity when possible
 2. survivable communication even when not connected
 3. technologies to enhance survivability
 - adaptive and agile networking
 - satellite and airborne nodes
 - end-to-end disruption tolerance
 - user controlled adaptive applications
- Summary

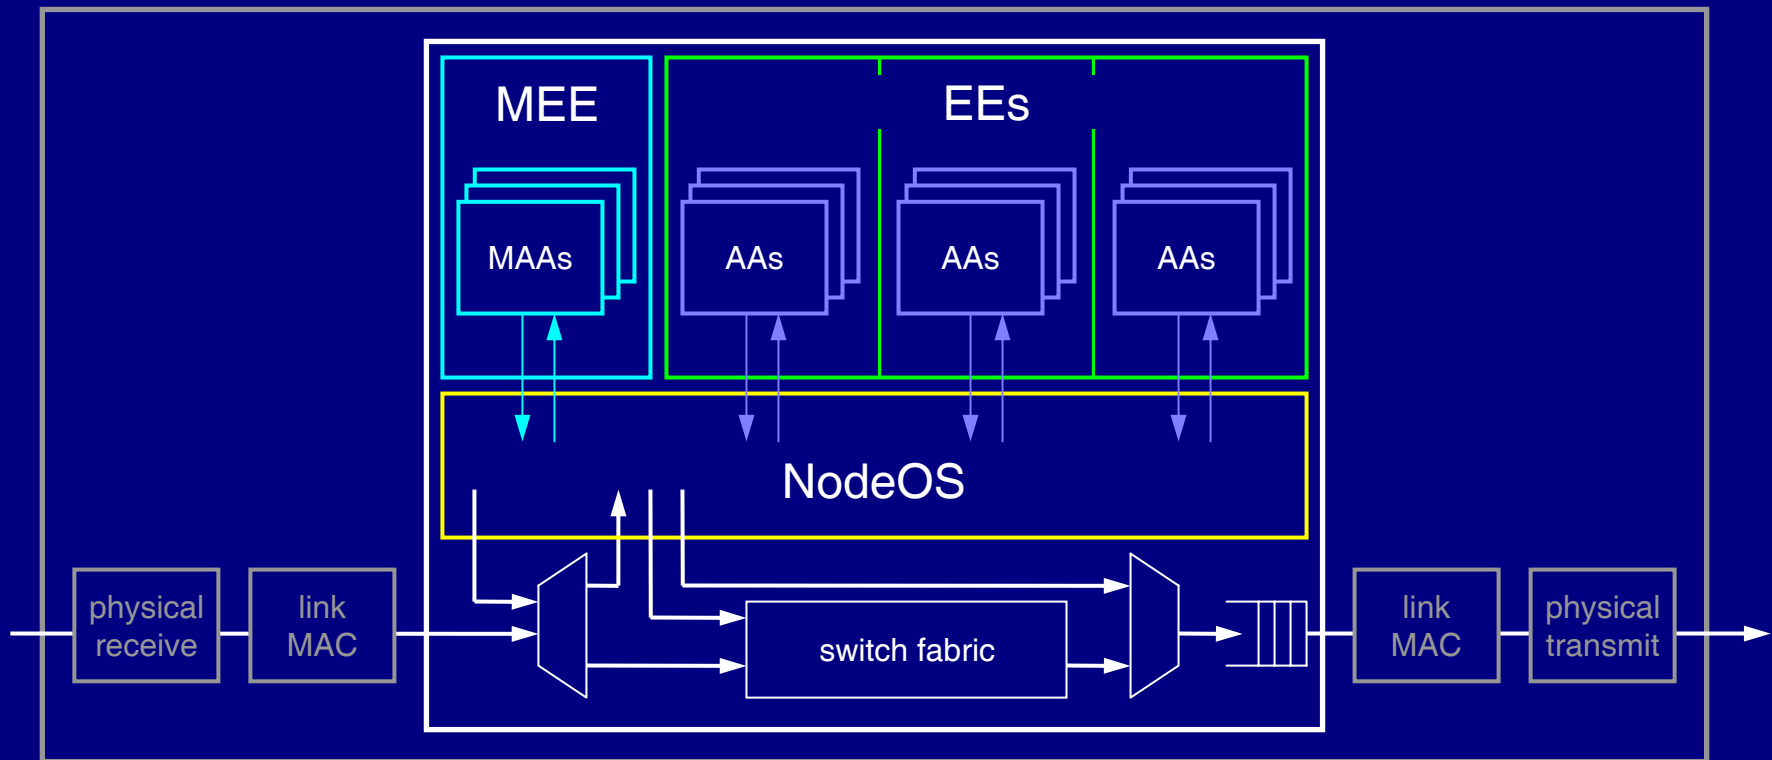
Survivability Technologies

Adaptive and Agile Networking

- Adaptive networking at all layers
 - physical layer transmission and coding
 - beamform, waveform, and frequency agility
 - MAC protocols and parameters
 - network routing, signalling, and addressing
 - geographical, topological, and characteristics-based
 - end-to-end protocols
- Enabled in systematic manner by
 - software radios, DSPs, and agile wideband RF transceivers
 - active networking technology
 - protocols and algorithms dynamically provisioned
 - don't need to standardise a network and MAC protocol
 - but rather a framework for negotiation

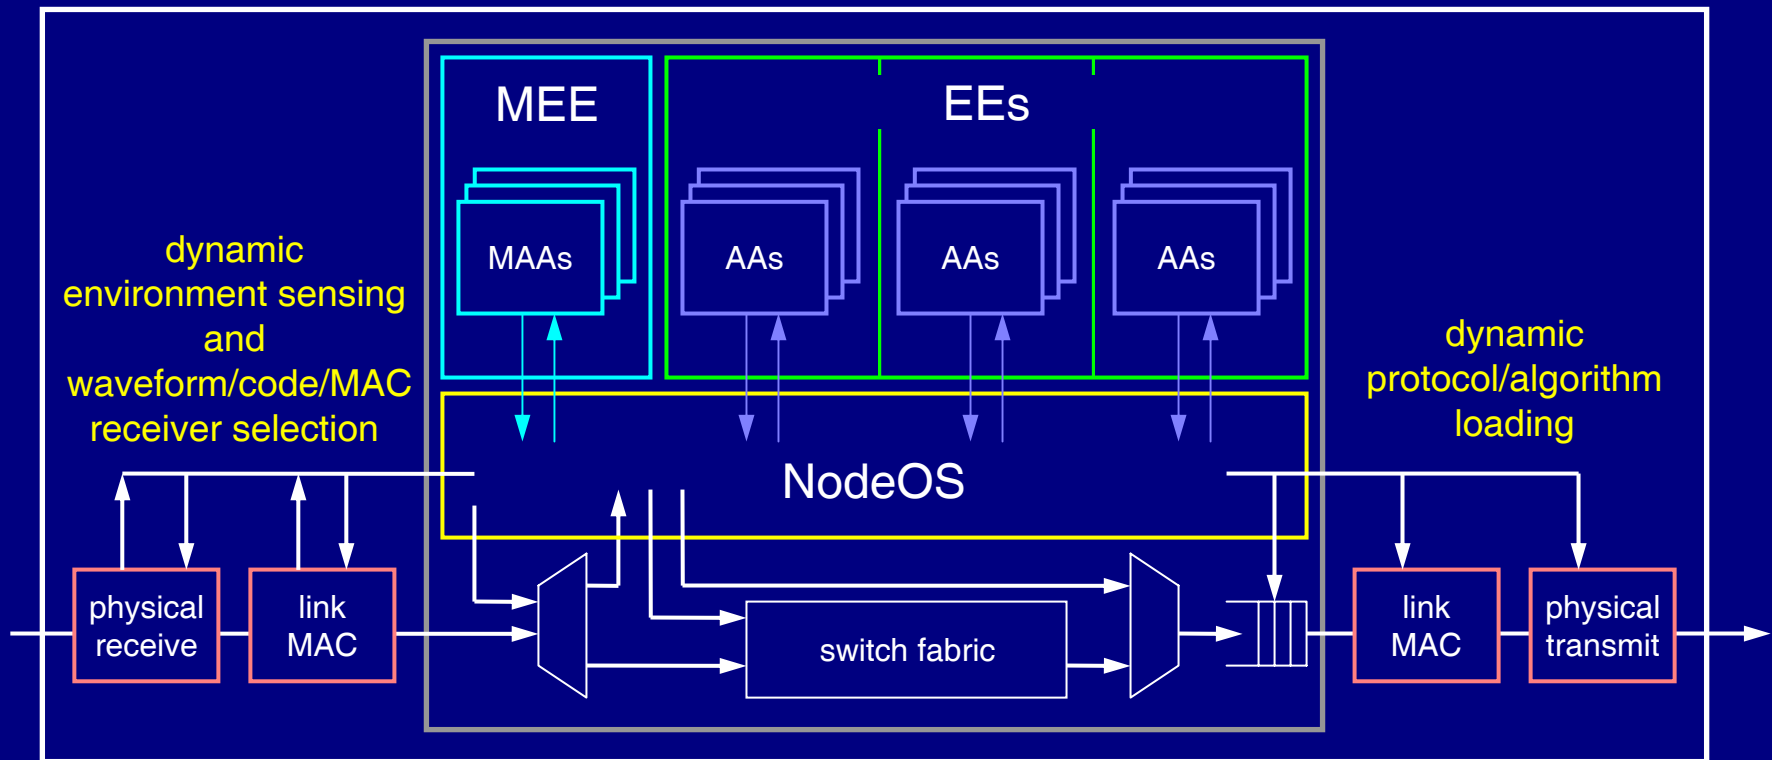
Survivability Technologies

DARPA Active Node Architecture



Survivability Technologies


Mobile Wireless Active Node Architecture



Survivability Technologies

Satellites and Airborne Nodes

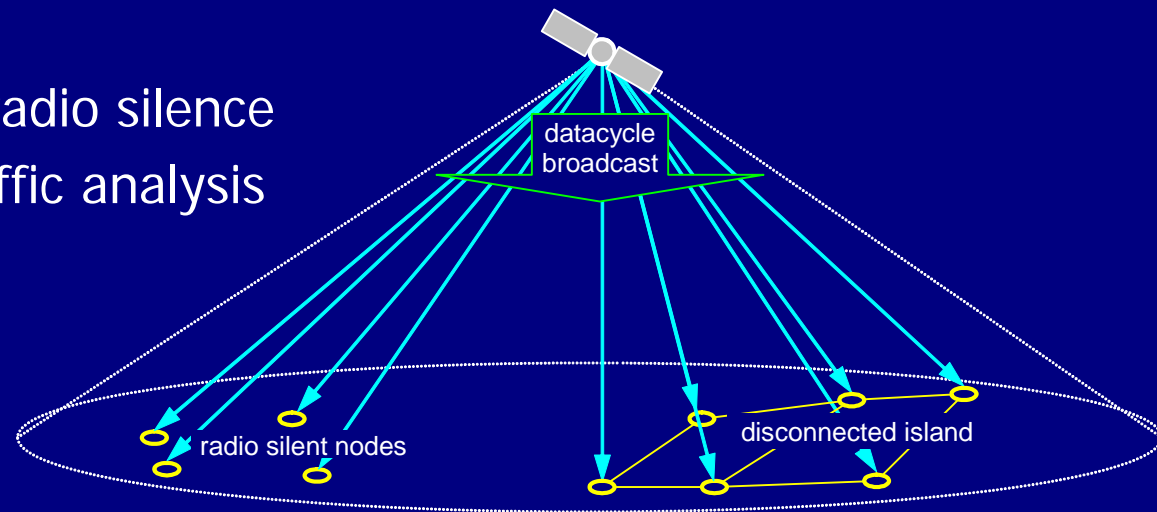
- High altitude vehicles (satellites and airborne nodes)
 - *less subject to line-of-sight obstructions*
 - less susceptible to attack
- Large transmission footprint
 - inherent broadcast
 - *connect disconnected islands*
 - inter-island relay
 - mitigates node mobility
 - transparent within footprint



Survivability Technologies

Airborne Information Dissemination

- Broadcast information dissemination
 - routing and topology updates
 - name services
 - certificates and CRLs
- Datacycle
 - supports radio silence
 - resists traffic analysis



Survivability Technologies

Satellites and Airborne Nodes

- Disadvantages
 - expensive and infrequent deployment
 - difficult to upgrade satellites
 - software upload
 - space shuttle mission for hardware & transmitter replacement
 - high transmission power required for uplink
 - mitigated by directional spot beam to known location
 - satellite
 - racetrack-path UAV (unpiloted aerial vehicle)

Survivability & Disruption Tolerance

End-to-End Survivability & Disruption Tolerance

- Introduction to survivability and disruption tolerance
 - definitions, characteristics
 - assumptions, requirements
- Survivability and disruption tolerance strategy
 - network survivability
 1. maintain survivable connectivity when possible
 2. survivable communication even when not connected
 3. technologies to enhance survivability
 - end-to-end survivability and disruption tolerance
 - disruption-tolerant user-controlled adaptive applications
- Summary


E2E Survivability and DT

Internet Transport Protocols

- Current transport protocols assume
 - strongly connected stable end-to-end paths
 - symmetric path
 - reliable medium
- TCP and SCTP
 - combined feedback error+flow+congestion control
 - reliable ACK stream required for self-clocking
 - unable to discriminate channel loss from congestion
 - congestion indicated by loss (switch drops)
 - channel loss results in throttling source
 - *wrong* response in uncongested network

E2E Survivability and DT

Knobs and Dials



- Transport protocols should interact with network
 - *dials* feed back topology and path characteristics
 - *knobs* influence lower layer parameters and behaviour
 - example: error control based on loss characteristics

E2E Survivability and DT


Asymmetric Paths

- Asymmetric channels result from
 - asymmetric transmission power
 - intentional (LPD) or available power
 - antenna characteristics and directionality
 - terrain and location
- Unidirectional channels result from
 - asymmetric transmission power
 - radio silence
- Path connectivity may be episodic
- Asymmetric and unidirectional E2E
 - concatenation of channels
 - forward and reverse may follow different paths



E2E Survivability and DT

Asymmetric End-to-End Paths

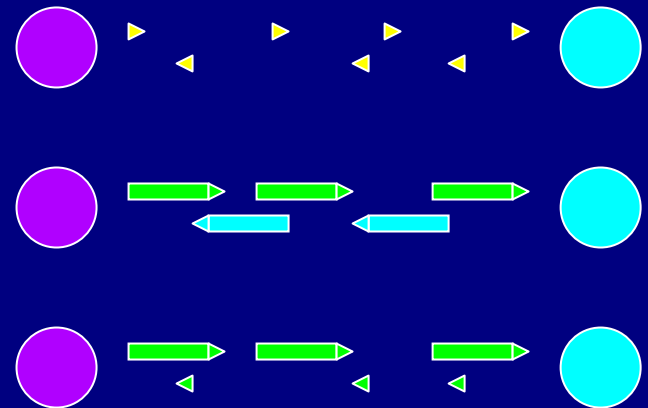


- Asymmetric end-to-end path challenges
 - how to find best paths through network
 - how to characterise entire path

E2E Survivability and DT

Bidirectional Paths


- Bidirectional path *required* for
 - pairwise synchronisation
 - signalling messages
 - bidirectional data communication
 - application issue
 - closed-loop feedback control
 - ACKs for *reliable* data transfer
 - *even* if data transfer unidirectional



E2E Survivability and DT

Open Loop Control

- Survivability with asymmetric channels needs:
 - open-loop control
 - with feedback only when necessary
- Open-loop rate control
 - congestion feedback from network only when necessary




E2E Survivability and DT

Open Loop Error Control

- Open-loop error control: **FEC**

- unreliable transfer
 - optional per link FEC
- quasi-reliable transfer
 - **FEC** for probabilistic reliability
- reliable transfer
 - requires bi-directional path
 - infrequent adaptive selective **ACKs**
 - distinct from:
 - flow control (E2E)
 - congestion control



note: SCTP does none of this

E2E Survivability and DT

End-to-End Transport Mechanisms

- Flow control
 - rate that *receiver* can accept
 - purely end-to-end
- Congestion control
 - rate that *network* can accept without congesting
 - network feedback to end systems
- Error control
 - retransmission of corrupt and lost packets
 - link and network-based error characteristics
 - application-dependent reliability requirements

E2E Survivability and DT

Explicit Loss/Congestion/Delay Discrimination

- Absence of expected packet or ACK arrival
 - three *distinct* and unrelated causes:
- 1. Congestion: packet dropped in network
 - congestion control: queue overflow (tail drop)
 - congestion avoidance: intentional packet drop
- 2. Corruption: packet lost or delivered corrupted
 - channel error causing bit errors
- 3. Delay: packet arrival later than expected
 - store-and-forward delays in disruption tolerant network
 - long path
 - speed-of-light delay in delay-tolerant network
 - very long path around disruption

E2E Survivability and DT

Discrimination and Explicit Notification

- Discrimination and proper response essential:
 - congestion \Rightarrow back off
 - corruption \Rightarrow retransmit
 - **delay** \Rightarrow **wait or**
retransmit via lower delay path
- Explicit notification
 - ECN: explicit congestion notification
 - ELN: explicit loss notification (due to corruption)
 - ELN *cannot* be determined from ECN (and vice versa)
 - packet that first causes congestion may then be corrupted

E2E Survivability and DT

Error Control Mechanisms

- Mechanisms
 - observation that error has occurred (detection)
 - notification of error
 - decision on what response to take
 - action to correct error
- Taxonomy of mechanisms
 - each mechanism may have different implementations...
...(e.g. E2E vs. HBH)...
 - ...but they are frequently related
- ETEN: explicit transport error notification


ETEN Taxonomy

Determinism and Granularity

- Determinism
 - deterministic (take action based on specific corruptions)
 - probabilistic (e.g. throttle source $x\%$ of the time)
- Granularity
 - PETEN: per packet response
 - CETEN: cumulative error rate response
- Control feedback ...
- Control locus ...
- Control band ...
- Control direction ...

ETEN Taxonomy

Control Feedback



- Closed loop feedback from notifier nodes $\{n, N\}$
 - (N)ACK from end system or switch (e.g. congestion drop)




- Open loop
 - Unreliable or FEC for statistical reliability, rate control




- Hybrid open + closed loop
 - FEC for statistical reliability + (N)ACKs as needed

ETEN Taxonomy


Control Locus



- End-to-end
 - no changes to network infrastructure




- Some hop-by-hop
 - deployed as needed (e.g. optional n at wireless access links)




- All hop-by-hop (N indicates *required* notifier functionality)
 - deployment challenges: impractical for Internet as a whole

ETEN Taxonomy

Control Band




- Out-of-band
 - ETEN signalling messages
 - may be forward or backward




- In-band
 - ETEN information carried by packets; options:
 - corrupted packets not dropped but marked (violates RFC 1812)
 - carried by subsequent packets in flow (header field or option)

ETEN Taxonomy

Control Direction



- Backward
 - ETEN messages returned by switches
 - out-of-band unless inserted in reverse flow
 - requires HBH ETEN



- Forward
 - ETEN forwarded to receiver; turned back to sender
 - may be in- or out-of band
 - longer delay in response for interactive over high bw- \times -delay

ETEN Simulation

TCP with Oracle ETEN

- TCP Oracle ETEN
 - instantaneous notification of corruption-based loss
 - perfect per packet loss discrimination
 - perfect response: retransmit rather than throttle
 - if loss within RTO
 - upper bound of possible benefit for bulk traffic
 - does *not* cause immediate retransmission:
would matter for transactions but not for bulk transfer
- Simulation
 - ns-2 **FullTcpSack**
 - 536B segments, 30 min. runs


ETEN Simulation

TCP with Oracle ETEN

LTN:
250ms 1-way
1.5Mb/s

single flow

30 min. run
536B segment
2400 seg window




ETEN Simulation

TCP with Oracle ETEN

LTN:
 250ms 1-way
 1.5Mb/s

competing:
 4xUDP
 CBR 0.25Mb/s
 on/off
 exp mean 0.5s

30 min. run
 536B segment
 2400 seg window



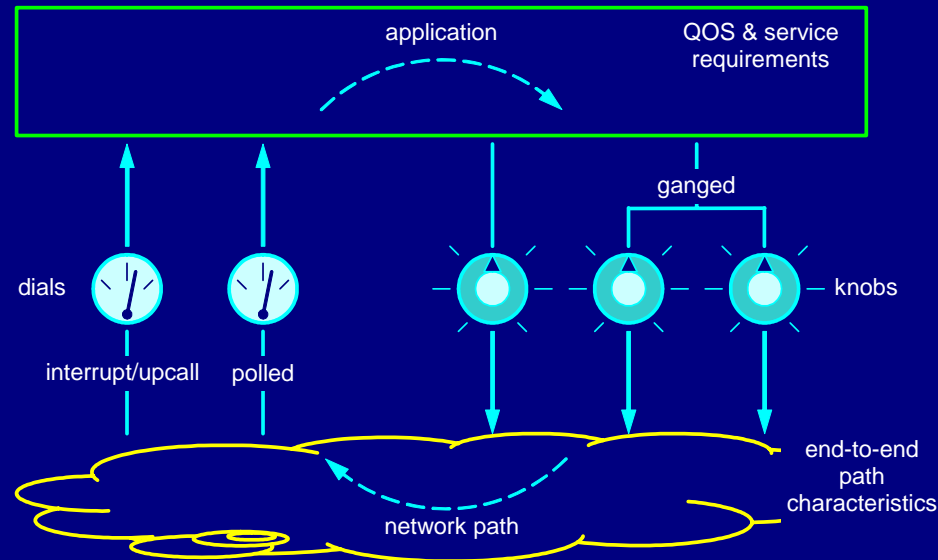
Survivability & Disruption Tolerance

User-Controlled Adaptive Applications

- Introduction to survivability and disruption tolerance
 - definitions, characteristics
 - assumptions, requirements
- Survivability and disruption tolerance strategy
 - network survivability
 1. maintain survivable connectivity when possible
 2. survivable communication even when not connected
 3. technologies to enhance survivability
 - end-to-end survivability and disruption tolerance
 - disruption-tolerant user-controlled adaptive applications
- Summary

Disruption Tolerant Applications

Adaptive with Knobs and Dials



- Applications should adapt to and mask disruptions
 - *dials* provide feedback from end-to-end paths (& network)
 - *knobs* influence transport and network behaviour
- Delay *translucency* (not transparency)

Disruption Tolerant Applications


Communication Association

- Recall: disruption tolerance goals for *applications*
 - information access by the user or application
 - end-to-end communication association
- Applications should *adapt* to path conditions
 - adaptive within modes
 - e.g. frame rate and resolution based on available bandwidth
 - adaptive between modes
 - e.g. video → audio → chat → messaging → email
- Driven by user preferences
 - e.g. tradeoff between frame rate and resolution

Disruption Tolerant Applications

Information Access

- Recall: disruption tolerance goals for *applications*
 - information access by the user or application
 - end-to-end communication association
- User experience highly dependent on response time
 - *interactive* information access
 - subsecond target response time
 - 100 ms ideal response time




Disruption Tolerant Applications

Example: WVM Latency-Aware Web Browser

- Distributed information access
 - access information from remote locations
 - Web provides most common infrastructure
 - web browser as client
 - HTTP as protocol
- WVM (Web VADE MECUM)
 - knobs and dials
 - user behaviour emulation
 - prototype: Mozilla on Linux using IBM Research WBI toolkit

Disruption Tolerant Applications

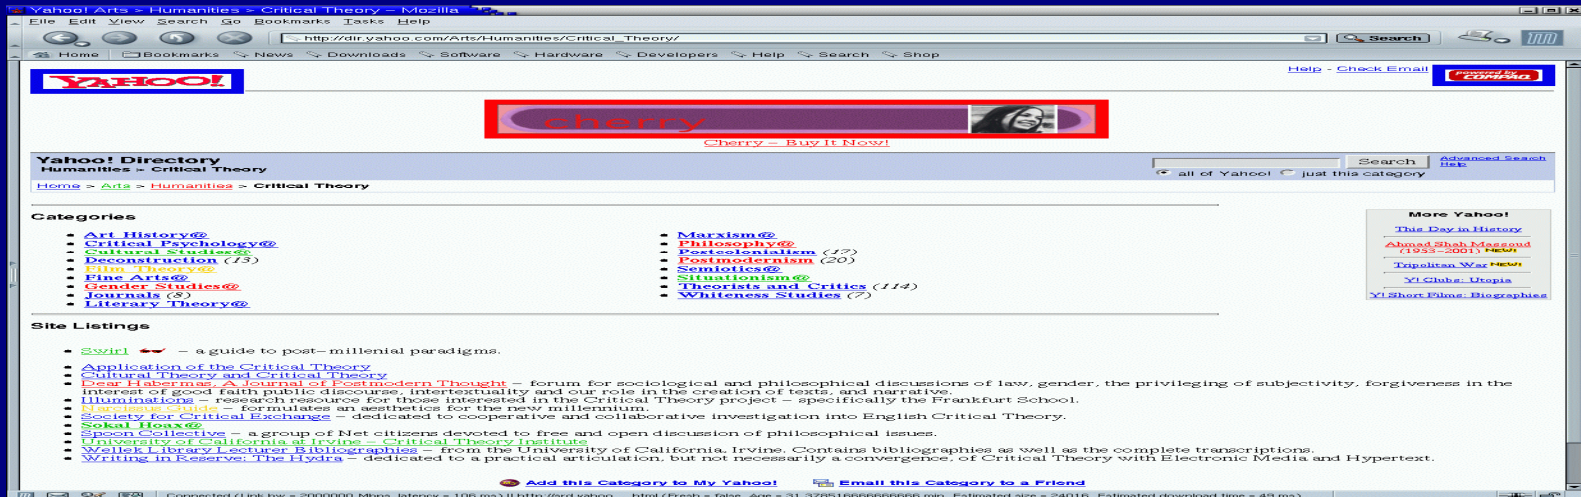
Example: WVM Knobs and Dials




- Dials
 - past response time t_r history*
 - weak link connectivity (instantaneous rate r_i to lighthouse)*
 - object size from server metadata
 - average end-to-end delay $l_i + l_s$ via probes to server
 - cached freshness* * implemented in prototype
- Knobs: application and user

Disruption Tolerant Applications

Example: WVM Dials – GUI



- Link color gives high level $\{fast, old, slow, unknown\}$

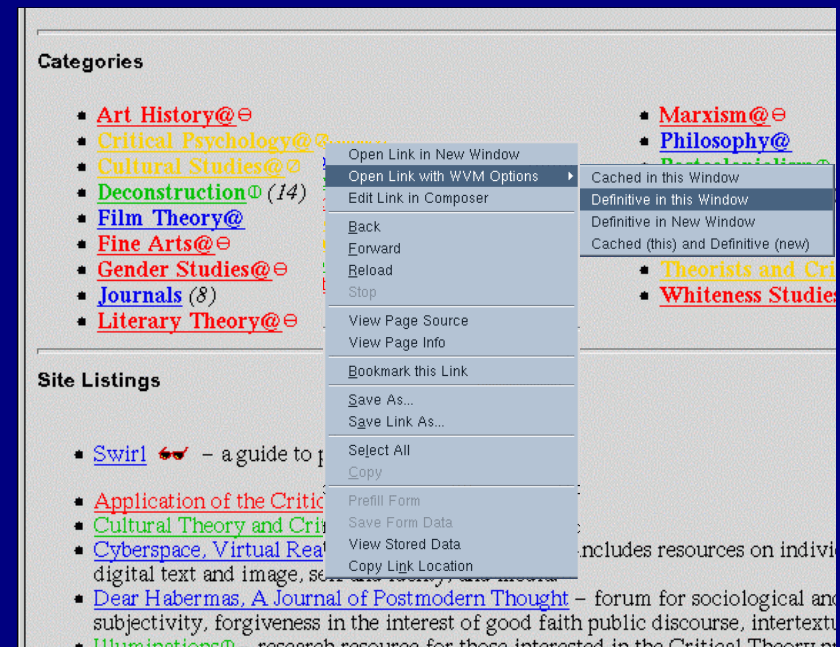


- Status bar indicates global and *per* link details
 - past response time
 - dynamic adjustment of weak link to lighthouse

Disruption Tolerant Applications

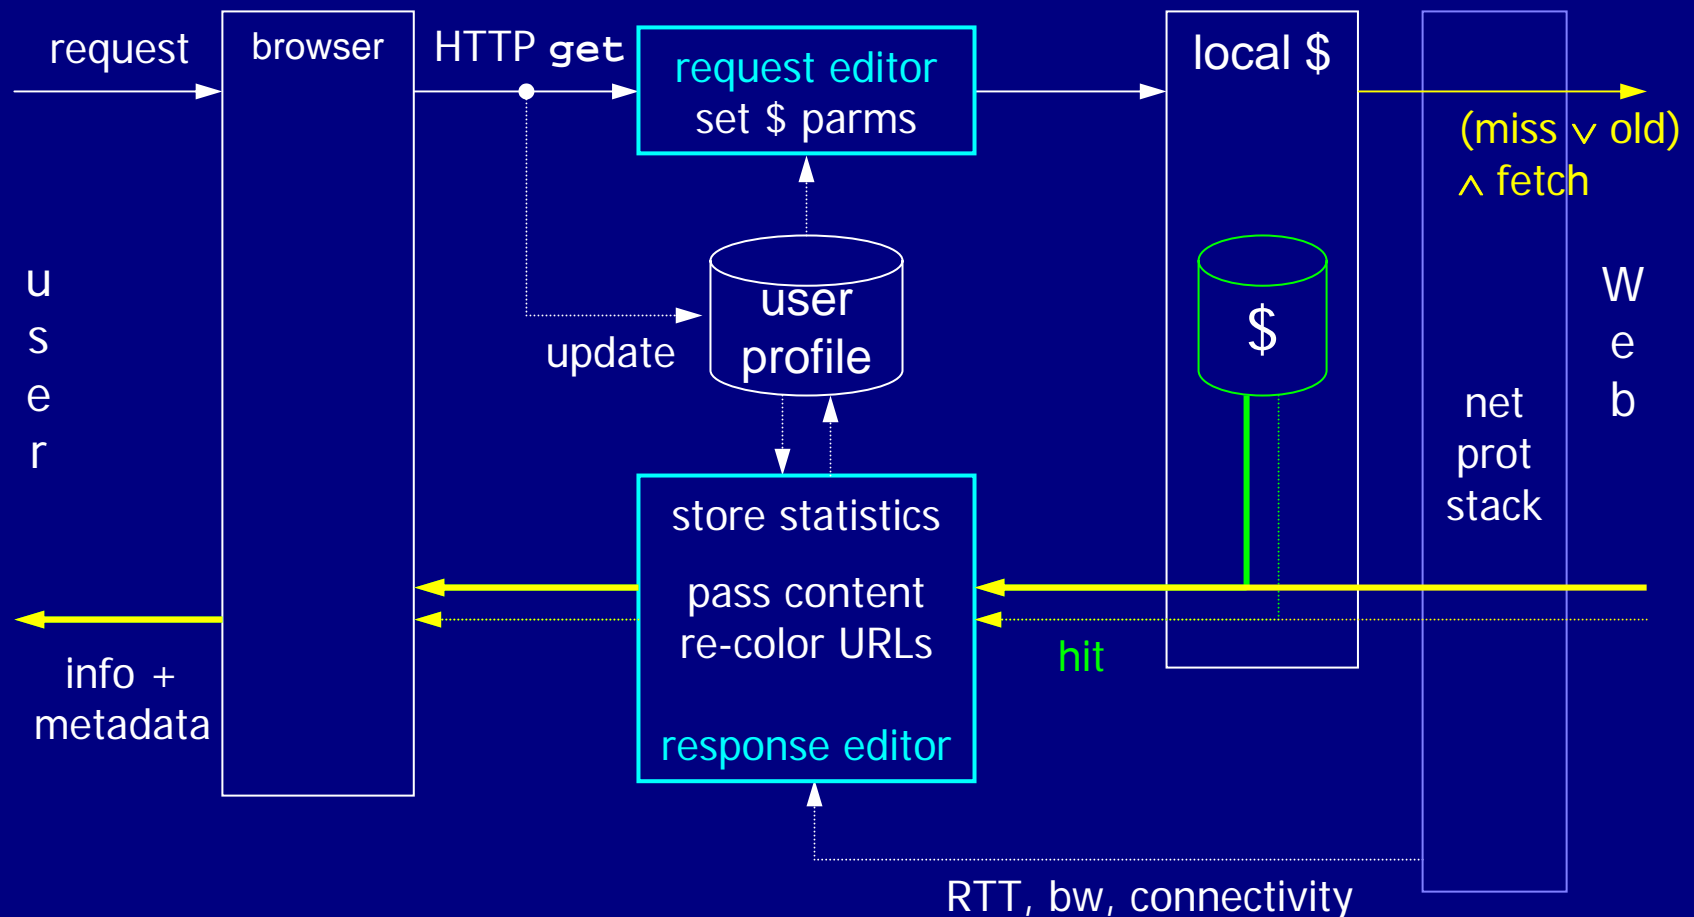
Example: WVM User Influence Knobs

- Fetch action
 - left click: default
 - get cached if available
 - profile based action
 - right click gives options:
 - fetch definitive
 - refresh window when definitive copy arrives
 - nonblocking fetch
 - definitive copy in new window when available
- View menu selection
 - allows display of unmodified page (un-munge HTML)



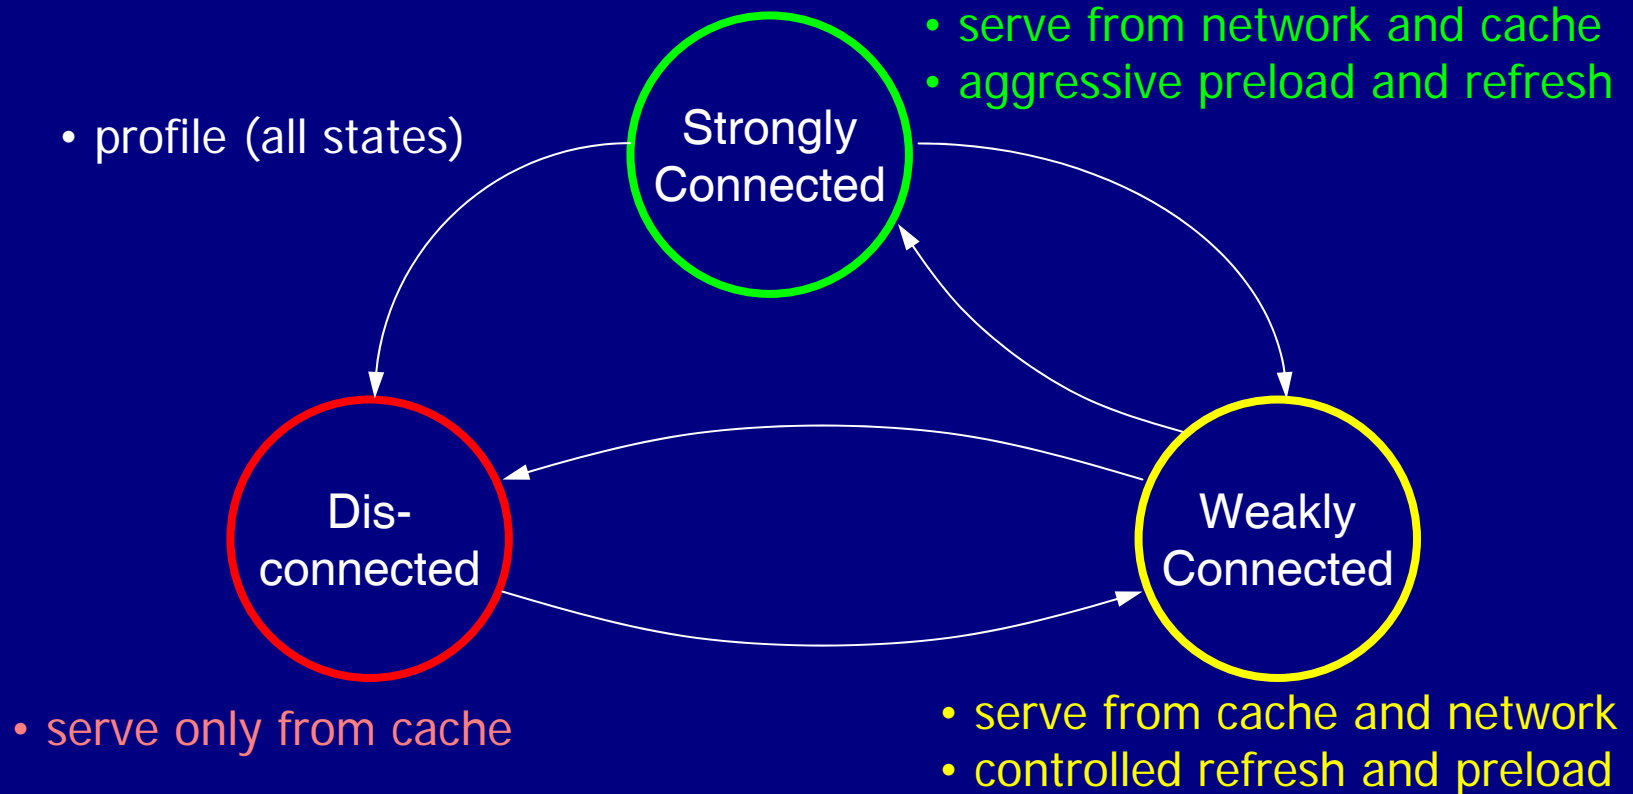
Disruption Tolerant Applications

Example: WVM Prototype Information Flow



Disruption Tolerant Applications

Example: WVM Proxy States



states based on CODA

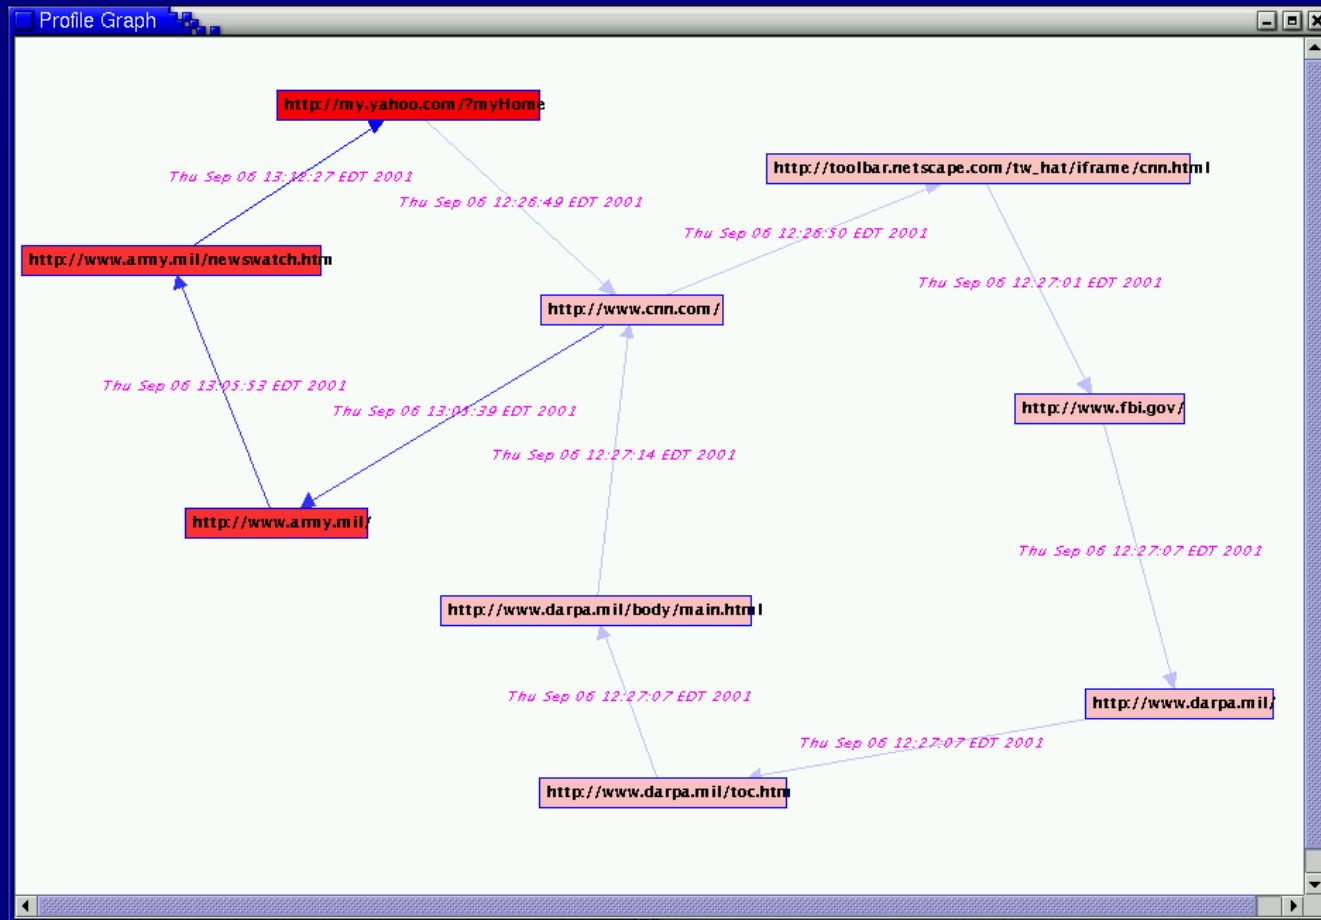
Disruption Tolerant Applications

Example: WVM Profile and User Emulation

- All user requests pass through WVM
- WVM builds a user-specific *WVM-ProfileGraph*
 - tracks and represents user Web access pattern
 - nodes are URLs accessed; directed edges for traversal
- URL priority limits cache life and refresh bandwidth
 - nodes and edges fade using a *decay formula*
- Autonomic user emulation
 - graph is automatically traversed when fully connected
 - graph slowly/selectively traversed when weakly connected
 - cache kept fresh even when user away from client
 - cache is as fresh as possible after state change {weak|disc}
 - coarse-grained learning possible (e.g. daily schedule)

Disruption Tolerant Applications

Example: WVM Profile Graph



Survivability & Disruption Tolerance Summary

- Introduction to survivability and disruption tolerance
 - definitions, characteristics
 - assumptions, requirements
- Survivability and disruption tolerance strategy
 - network survivability
 1. maintain survivable connectivity when possible
 2. survivable communication even when not connected
 3. technologies to enhance survivability
 - end-to-end survivability and disruption tolerance
 - disruption-tolerant user-controlled adaptive applications
- Summary

Survivability & Disruption Tolerance

Summary

- Attack problem at *all* levels:
 - physical, MAC, and link robustness and agility
 - network survivability
 - end-to-end survivability and disruption tolerance
 - disruption-tolerant user-controlled adaptive applications
- Beyond fault tolerance and crypto
- Design for survivability and disruption tolerance
 - expect challenging communication channel environment
 - expect and exploit mobility
 - expect, adapt, and mask high latency with user influence
 - interlayer awareness and control (knobs and dials)
 - intelligent resource & constraint tradeoffs (P, M, B, E, L)

Survivability & Disruption Tolerance

Primary References

Available from <http://www.sterbenz.org/sumowin>

SUMOWIN/DTN

James P.G. Sterbenz, Rajesh Krishnan, Regina Rosales Hain, Alden W. Jackson, David Levin, Ram Ramanathan, and John Zao, “Survivable Mobile Wireless Networks: Issues, Challenges, and Research Directions”, *Proceedings of the Wireless Security Workshop (WiSE) 2002*, MobiCom Atlanta, GA, Sep. 2002, pp.31–40.

ETEN

Rajesh Krishnan, James P.G. Sterbenz, Wesley M. Eddy, Craig Partridge, and Mark Allman, “Explicit Transport Error Notification (ETEN) for Error-Prone Wireless and Satellite Networks”, *Computer Networks*, vol.46 #3, Oct. 2004, pp. 343–362.

Rajesh Krishnan, Mark Allman, Craig Partridge, and James P.G. Sterbenz, *Explicit Transport Error Notification (ETEN) for Error-Prone Wireless and Satellite Networks*, BBN Technical Report 8333, Feb. 2002.

WVM

James P.G. Sterbenz, Tushar Saxena, and Rajesh Krishnan, *Latency-Aware Information Access with User-Directed Fetch Behaviour for Weakly-Connected Mobile Wireless Clients*, BBN Technical Report 8340, May. 2002.

Long latency

James P.G. Sterbenz and Joseph D. Touch, *High Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication*, John Wiley Networking Council (A. Lyman Chapin technical editor), Apr. 2001

Survivability & Disruption Tolerance Acknowledgements

- DARPA
 - Doug Maughan* (SUMMOWIN)
 - Rob Ruth* (GloMo)
 - Jean Sholtz* (WiaB)
- NASA
 - William Ivancic
- UMass
 - Jim Kurose
- Lancaster University
 - David Hutchison
- BBN
 - Mark Allman*
 - Alden W. Jackson
 - Ram Ramanathan
 - Tushar Saxena*
 - Martha Steenstrup*
 - Fabrice Tchakountio

*some of this work performed at
BBN Technologies*

* former affiliation

Survivability & Disruption Tolerance

Recent Presentation Venues

- 28 Apr 2004 Boston University
- 18 Jun 2004 KTH, Stockholm, Sweden
- 1 Oct 2004 Universität Tübingen, Germany
- 2 Dec 2004 Lancaster University, UK
- 16 Dec 2004 University of Kansas (abridged)
- 7 Feb 2005 Lancaster University, UK

End of Foils