Content-Centric Networking

J.J. Garcia-Luna-Aceves

UCSC and PARC

jj@soe.ucsc.edu jjgla@parc.com

Outline

- Problem we address
- Limitations of routing schemes that assume connected networks
- Our progress and initial steps
- Next steps and future direction

IP Internet Today

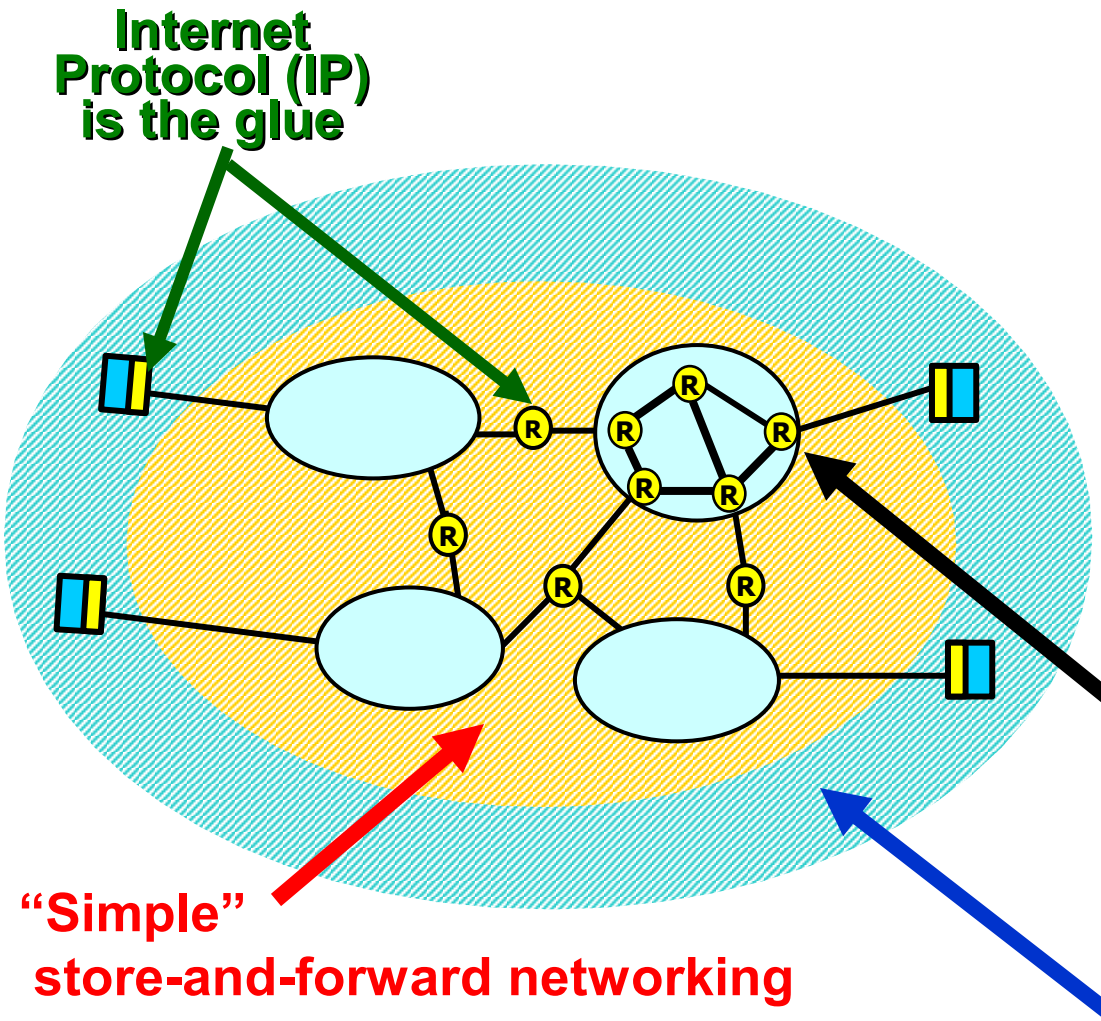

A Success tale of "two worlds with a little glue"

"Networking" is independent of processing and storage of content.

How Can We Live with Disruption?

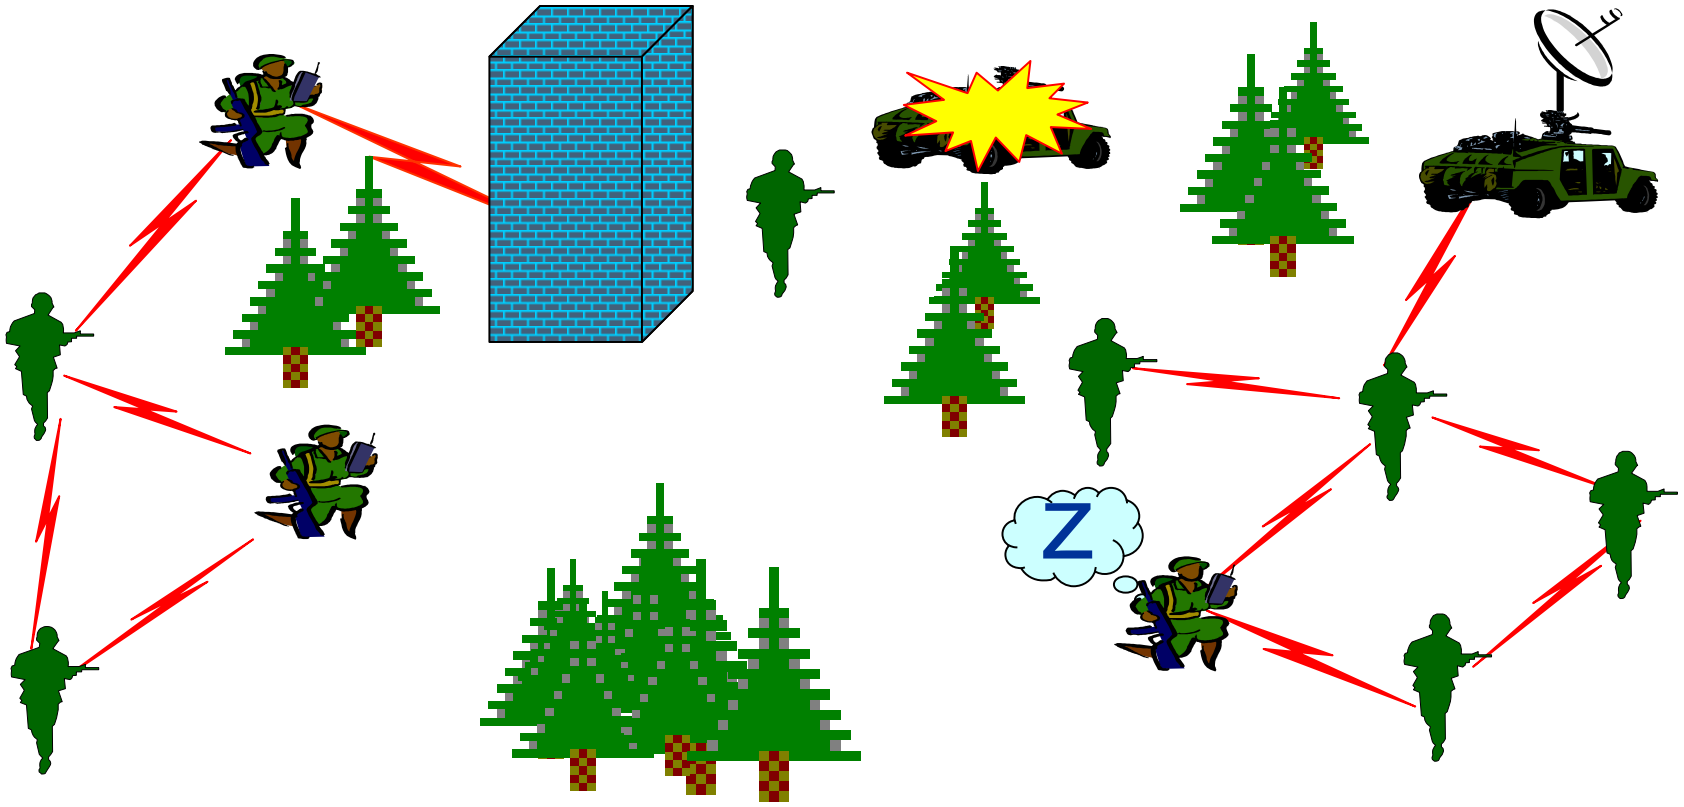

End-to-end connectivity *need not ever exist* and links (contacts) may not be suitable for schedules

Use Storage, Processing, and Communication Opportunistically

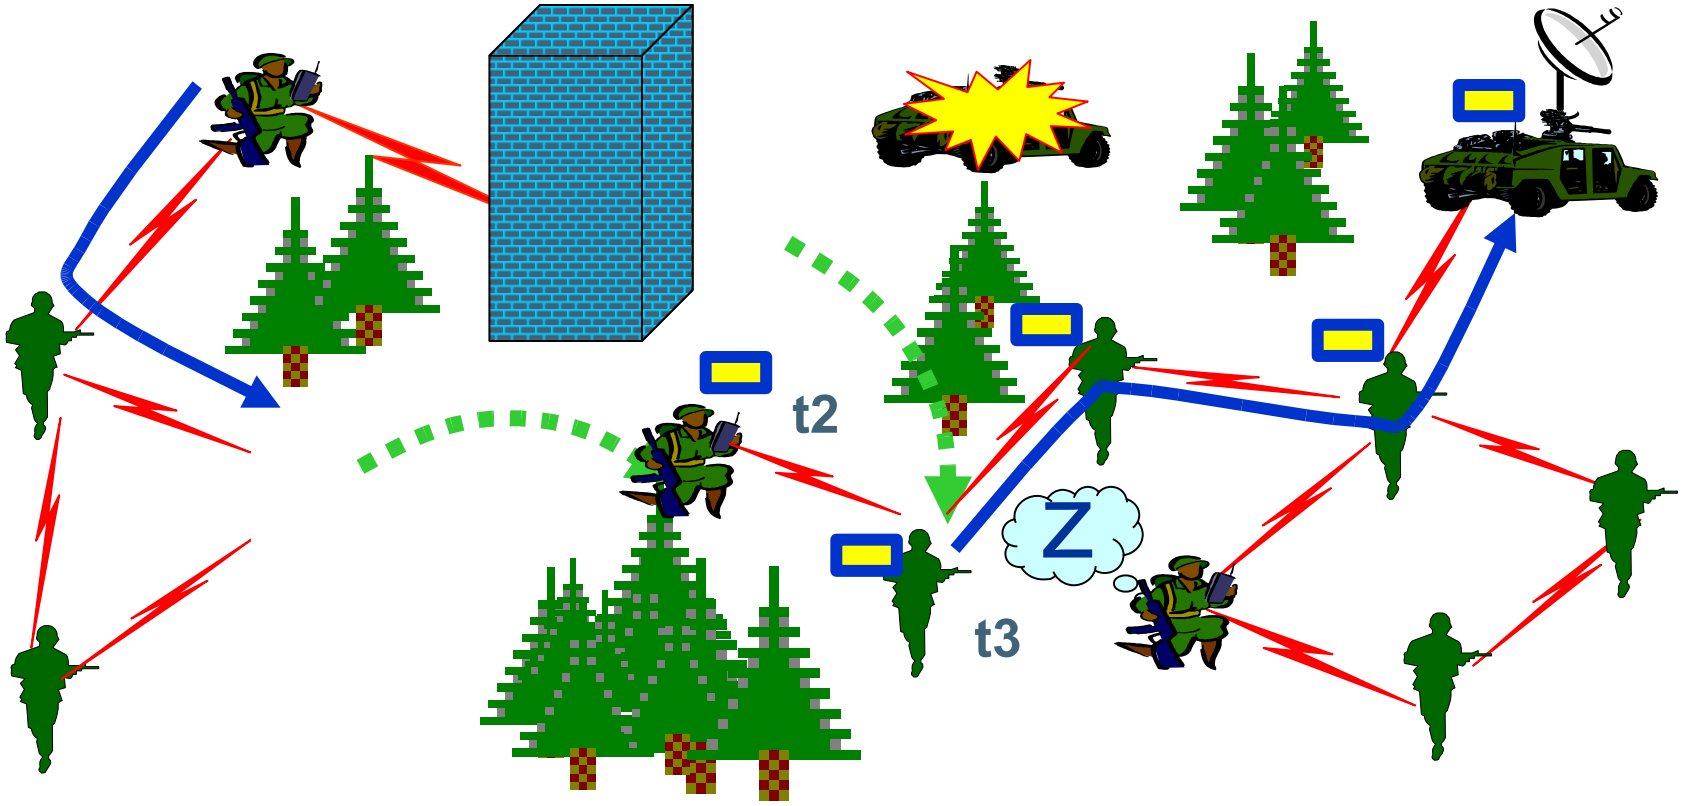

Treat *routes as functions of space and time*

Exploit longer-term storage of nodes

Opportunistic "store-process-forward"

Limitations of Prior Routing Approaches

- Routing independent of time-dependency of links:
  - Proactive routing
  - On-demand routing
  - Epidemic routing
- Routing that considers space-time constraints of links (contacts) works if we can assume the ability to know schedules of links (Oracles)

Proactive Routing:

Too many nodes are forced to know about how to reach each destination! Does not work well with random partitions

Path first, then data forwarding

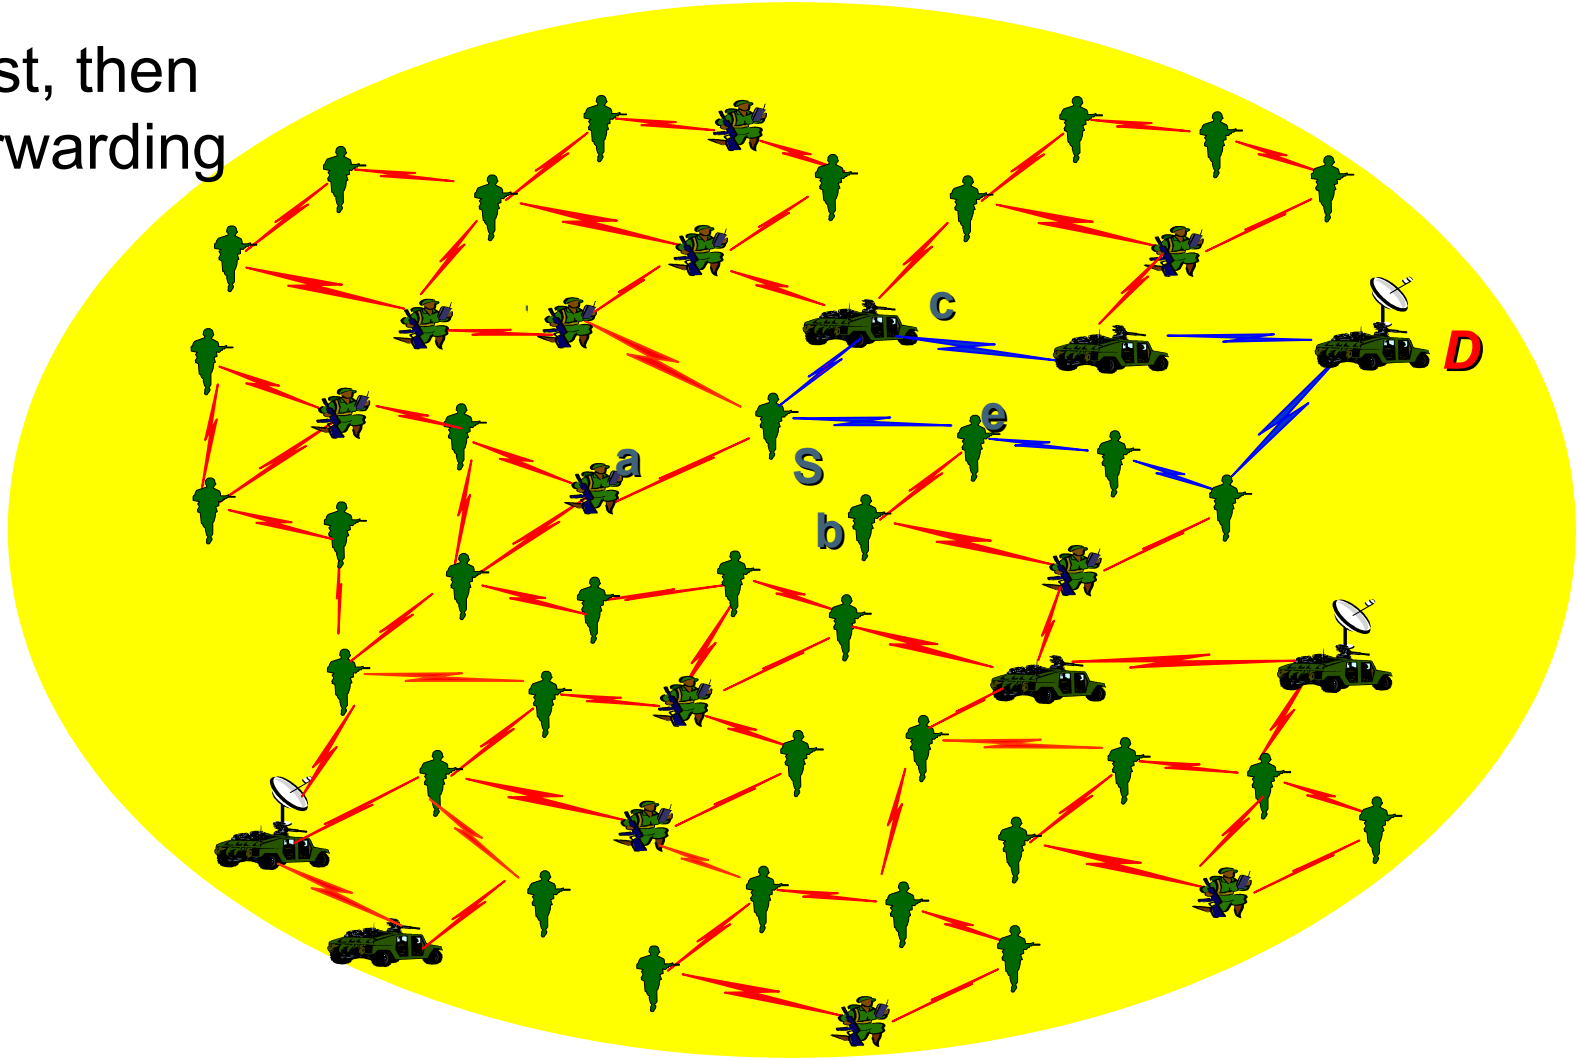

On-Demand Routing:

Too many nodes are forced to help find or repair ways to reach a few destinations! (RREQ flooding). Does not work with partitioned networks!

Path first, then data forwarding

Nodes with paths to D reply to S.

Too few nodes keep state for D. So too many nodes try to fix broken paths

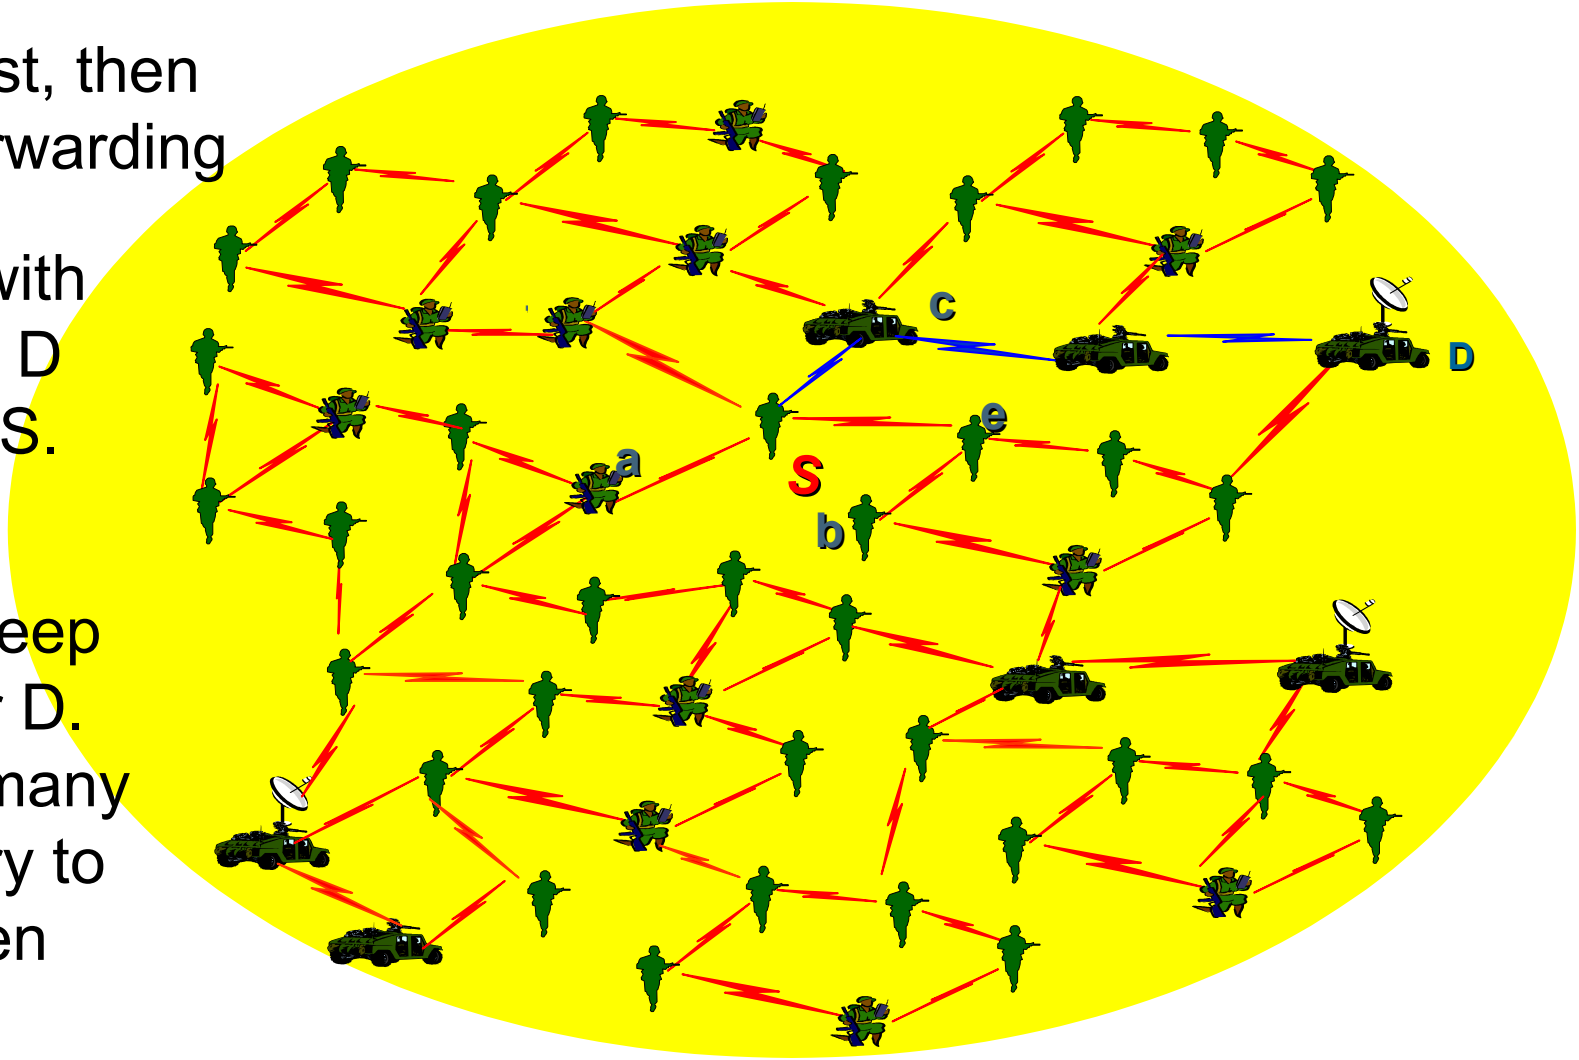

Epidemic Routing

Too many nodes are forced to relay data from S to D.
Does not work with partitioned networks, unless infinite storage is assumed.

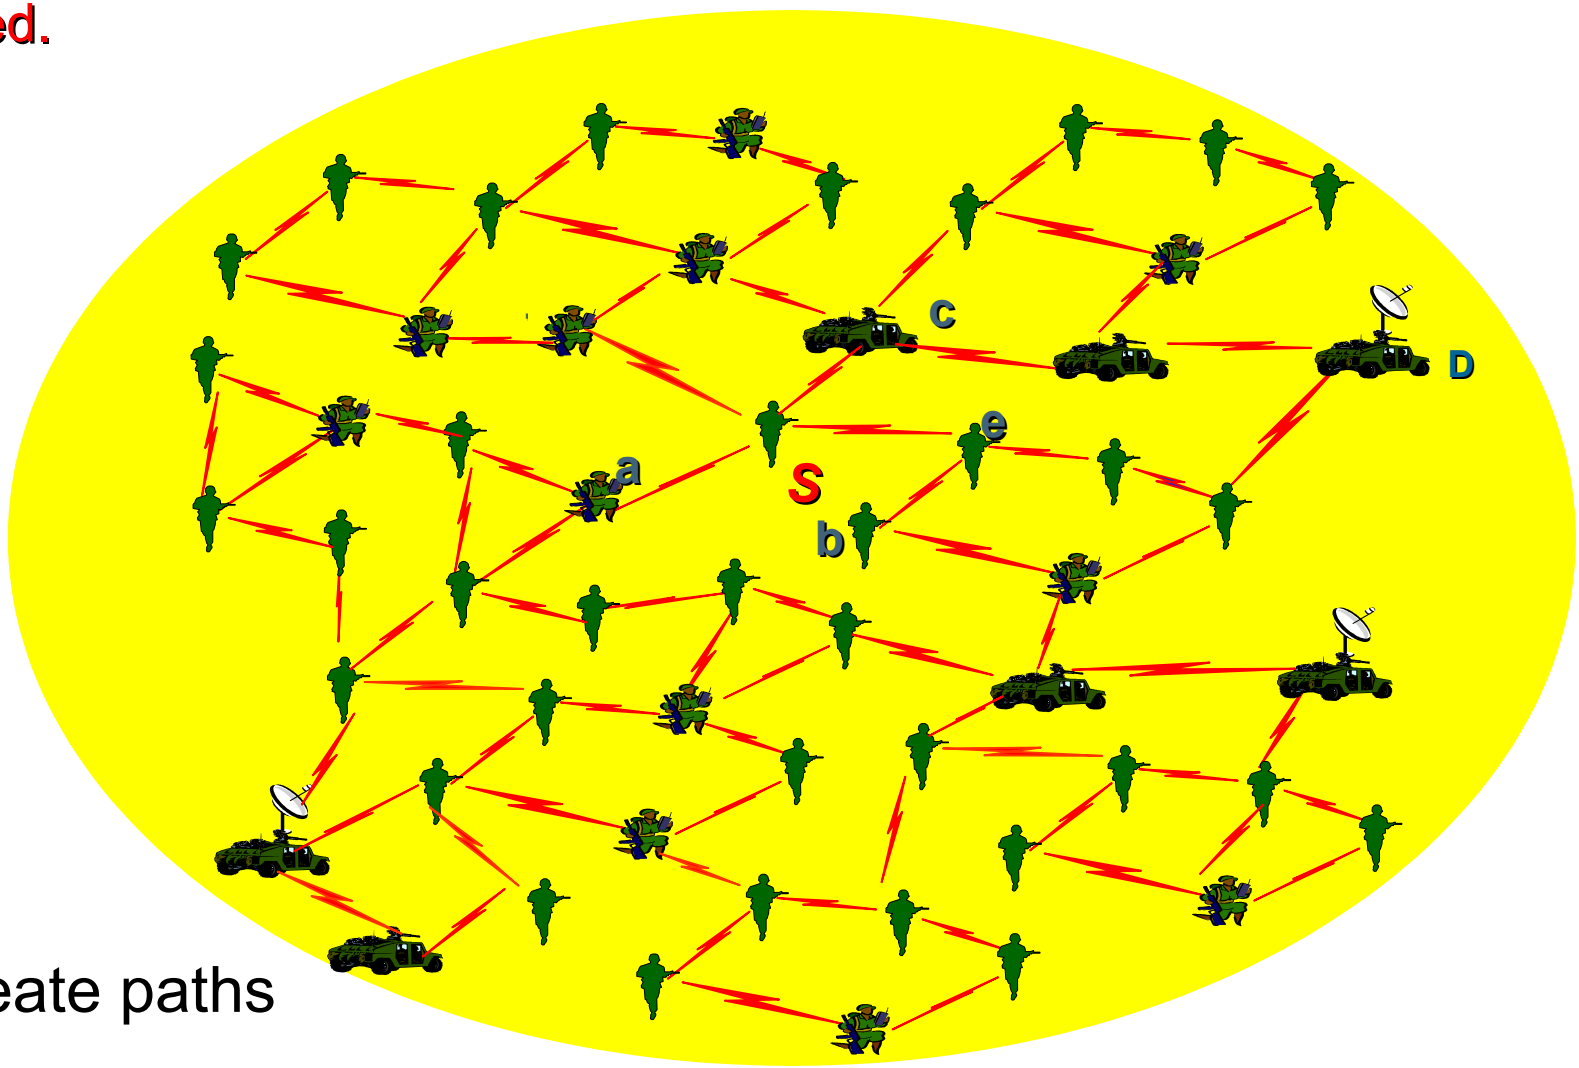

Data create paths

Goals

Limit the number of nodes that incur signaling and forwarding overhead between S and D

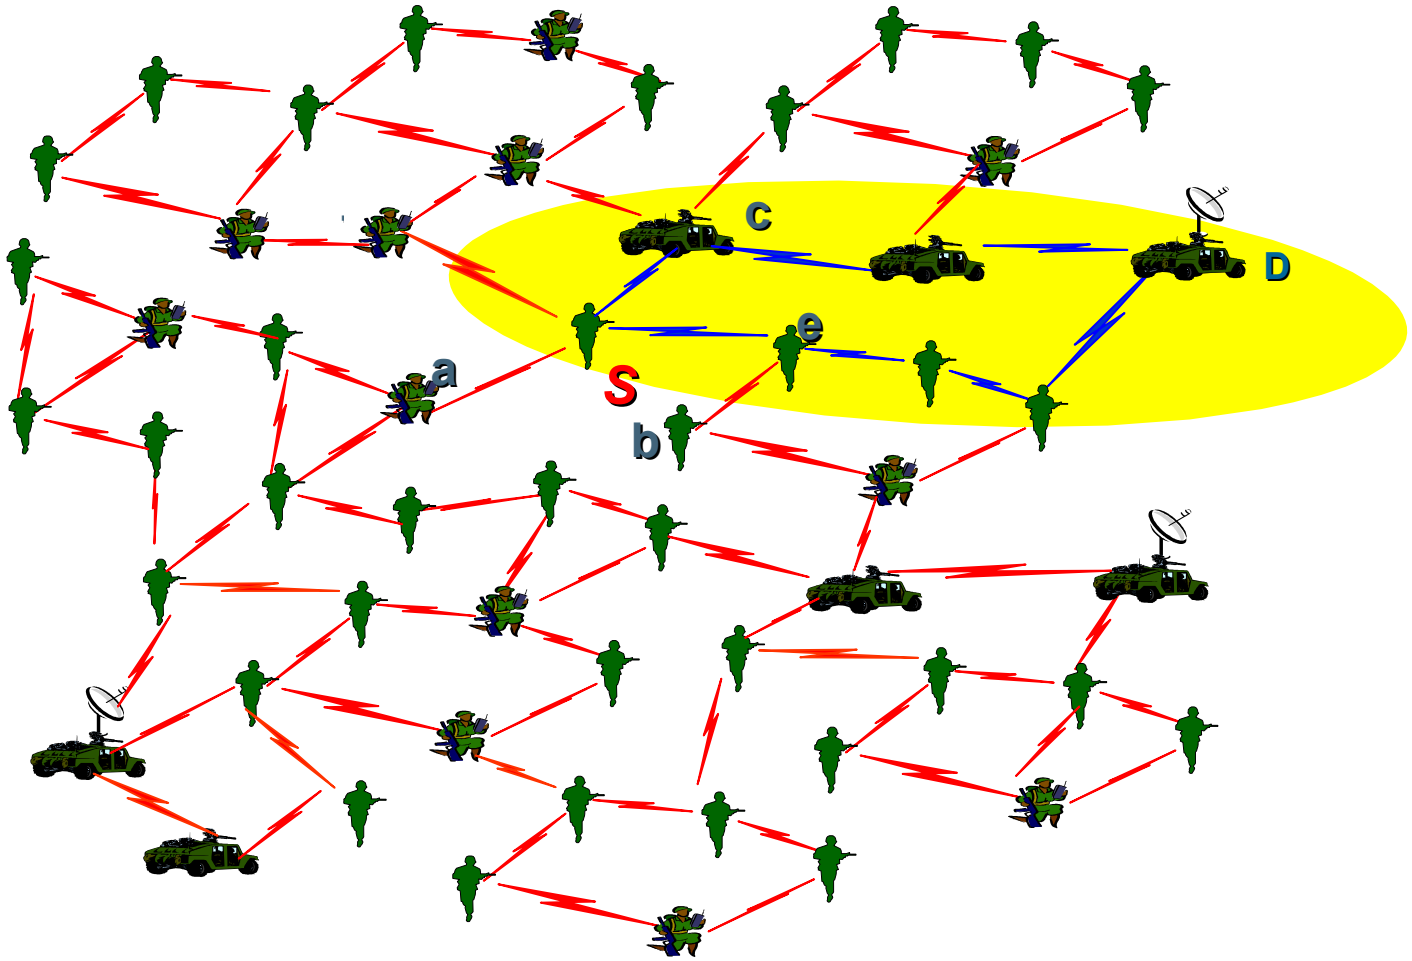

Goals

Enable Correct Signaling and Forwarding in Partitioned Networks. Preserve efficiency in each network component

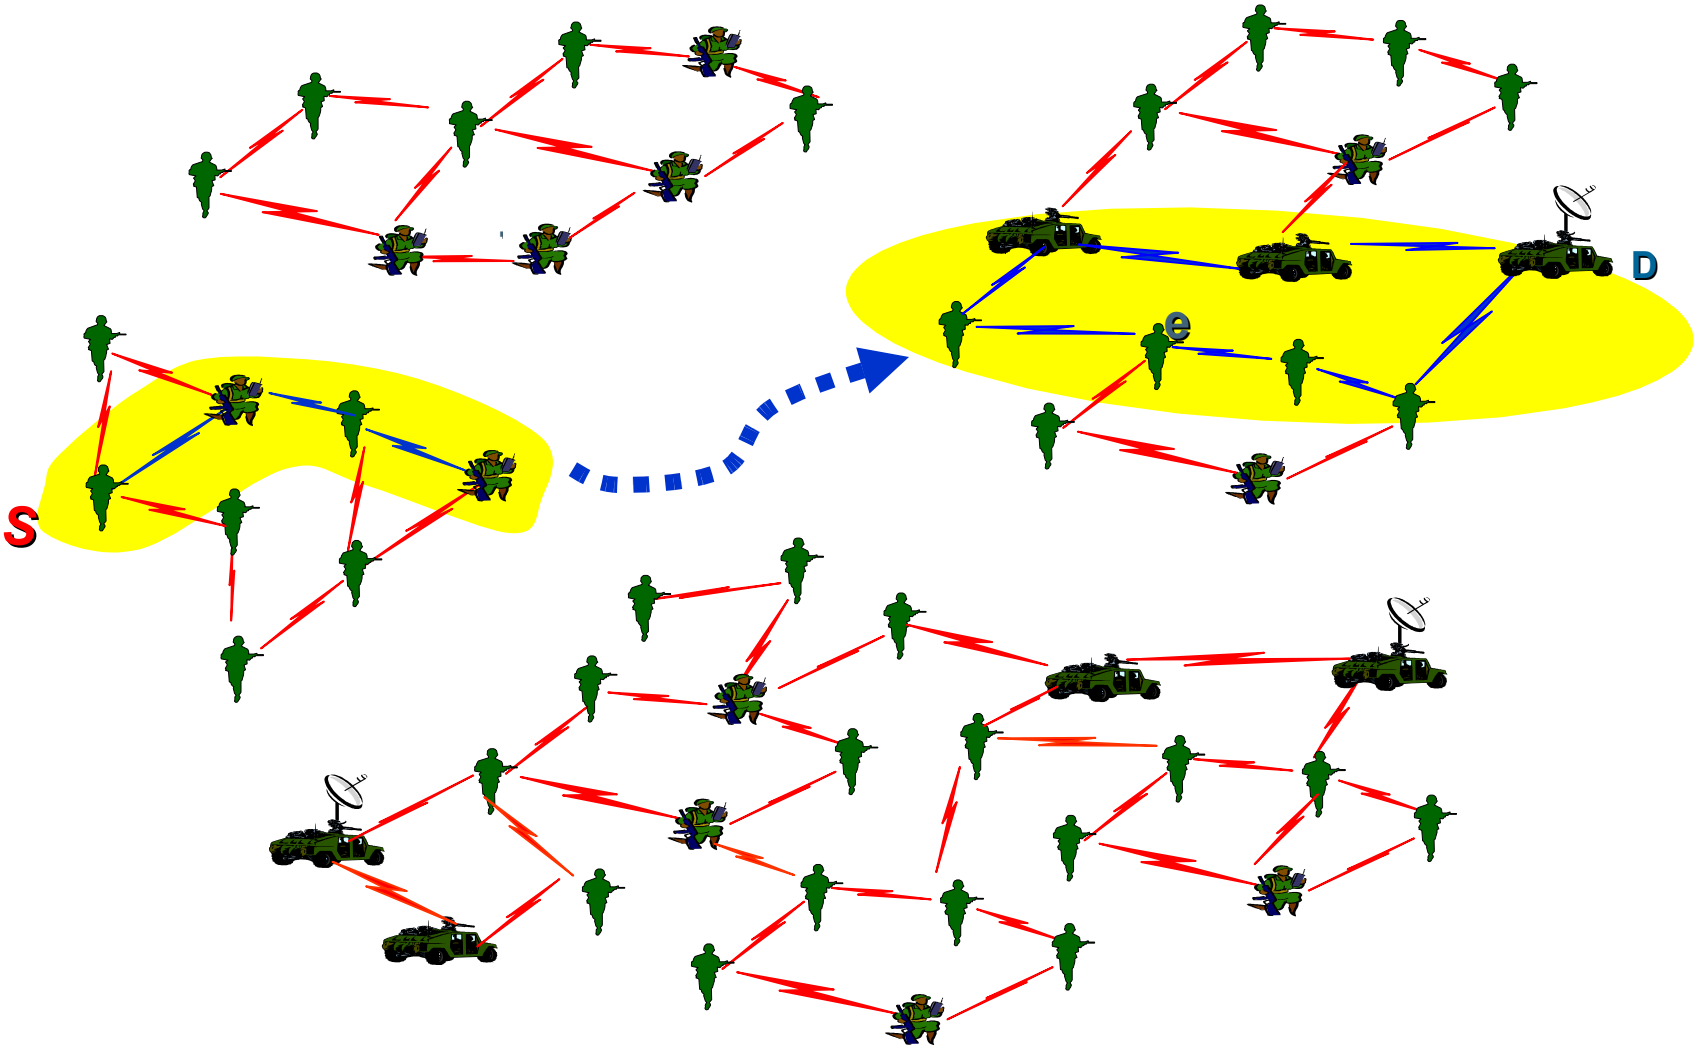

Steward Assisted Routing (StAR)

- ***SCIP (scoped contact and interest propagation):***
  - Destinations of interest are found with “interest messages” (like RREQs) stating the destination and duration of interest.
  - *Content* (data and signaling) states how long it needs to live!
  - Once found, destinations of interest (and stewards) start advertising themselves proactively
  - Advertisements propagate within the horizon of the “most distant interest”.
- ***Stewards***
  - Those nodes who are most likely to deliver a message to its intended destination (use last seq # heard from D and hops traversed by seq #).
  - Elected within each component for which destination of interest is known [by most recent (transitive) contact with the destination].
  - Loop-free routes maintained to stewards within a component, and among stewards towards destination across components [right now using destination & steward seq. #]

Example

S floods its interest in D within its connected component network component with some lifetime

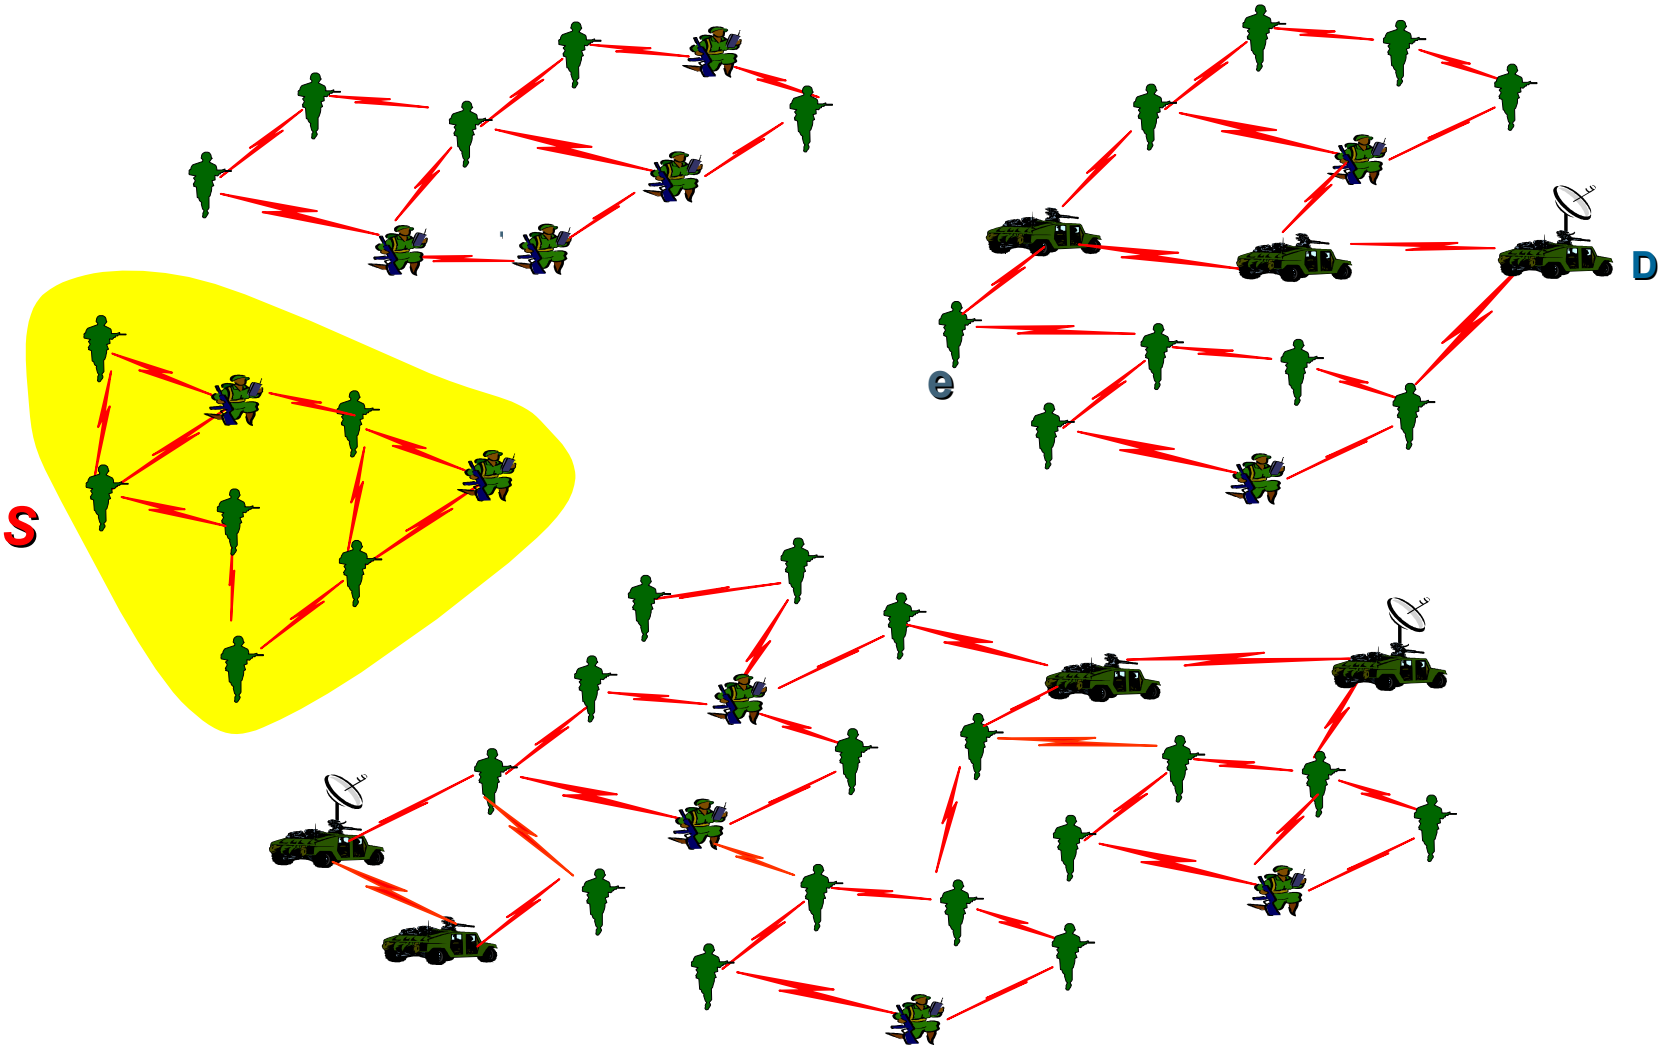

Example

Node f moves close to component and hears interest.
Assume e already knows about D.

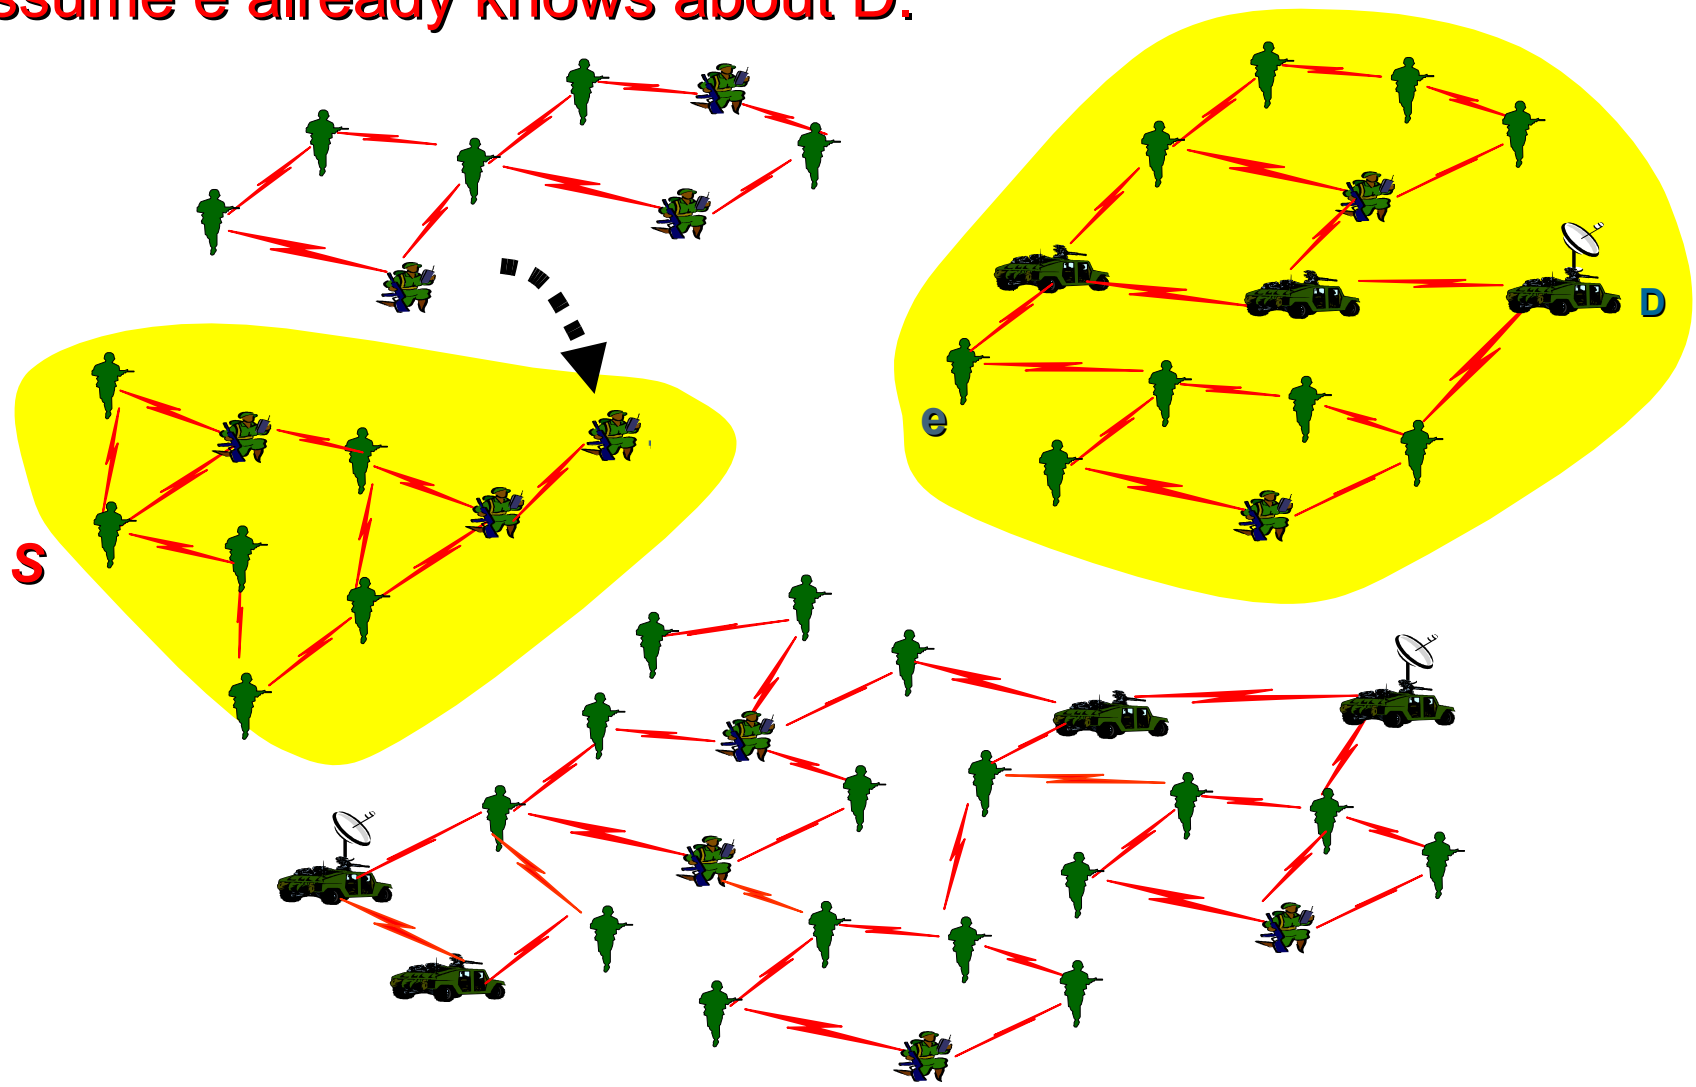

Example

Node f moves close to e, who conveys a seq # for D with 3 traversed hops.

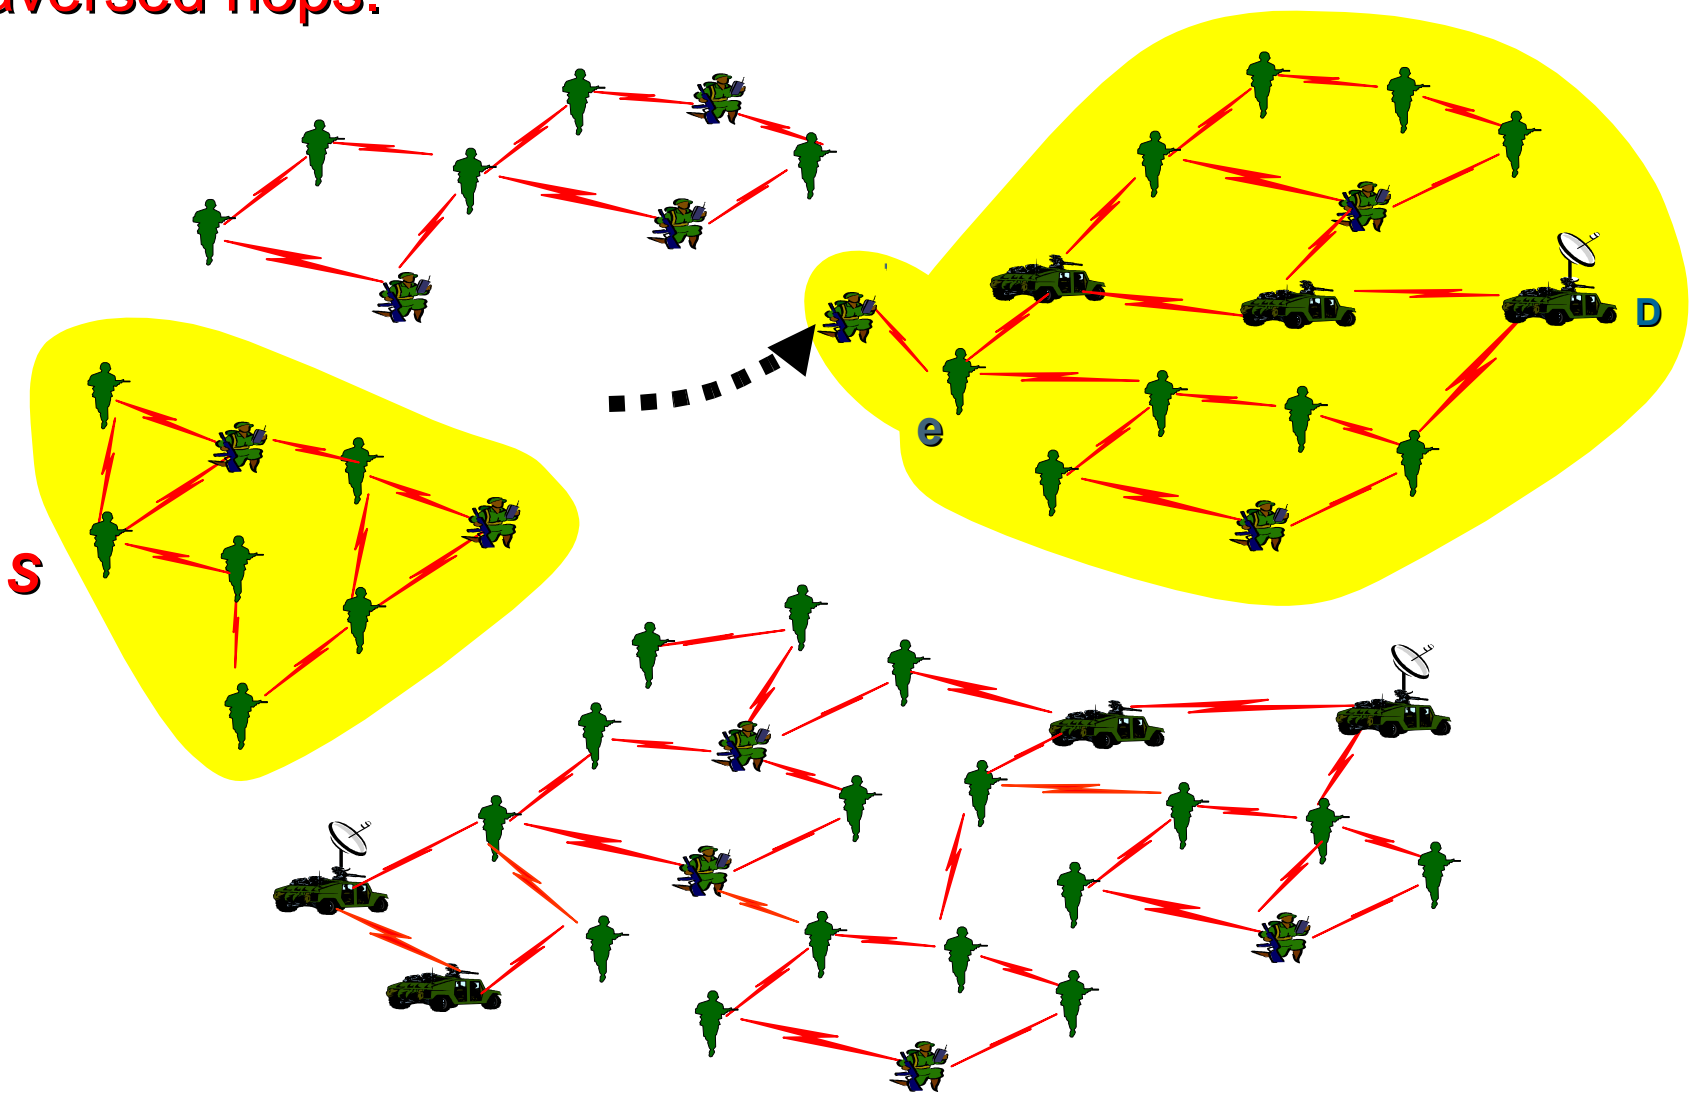

Example

Node f moves back close to S and becomes steward in the component.

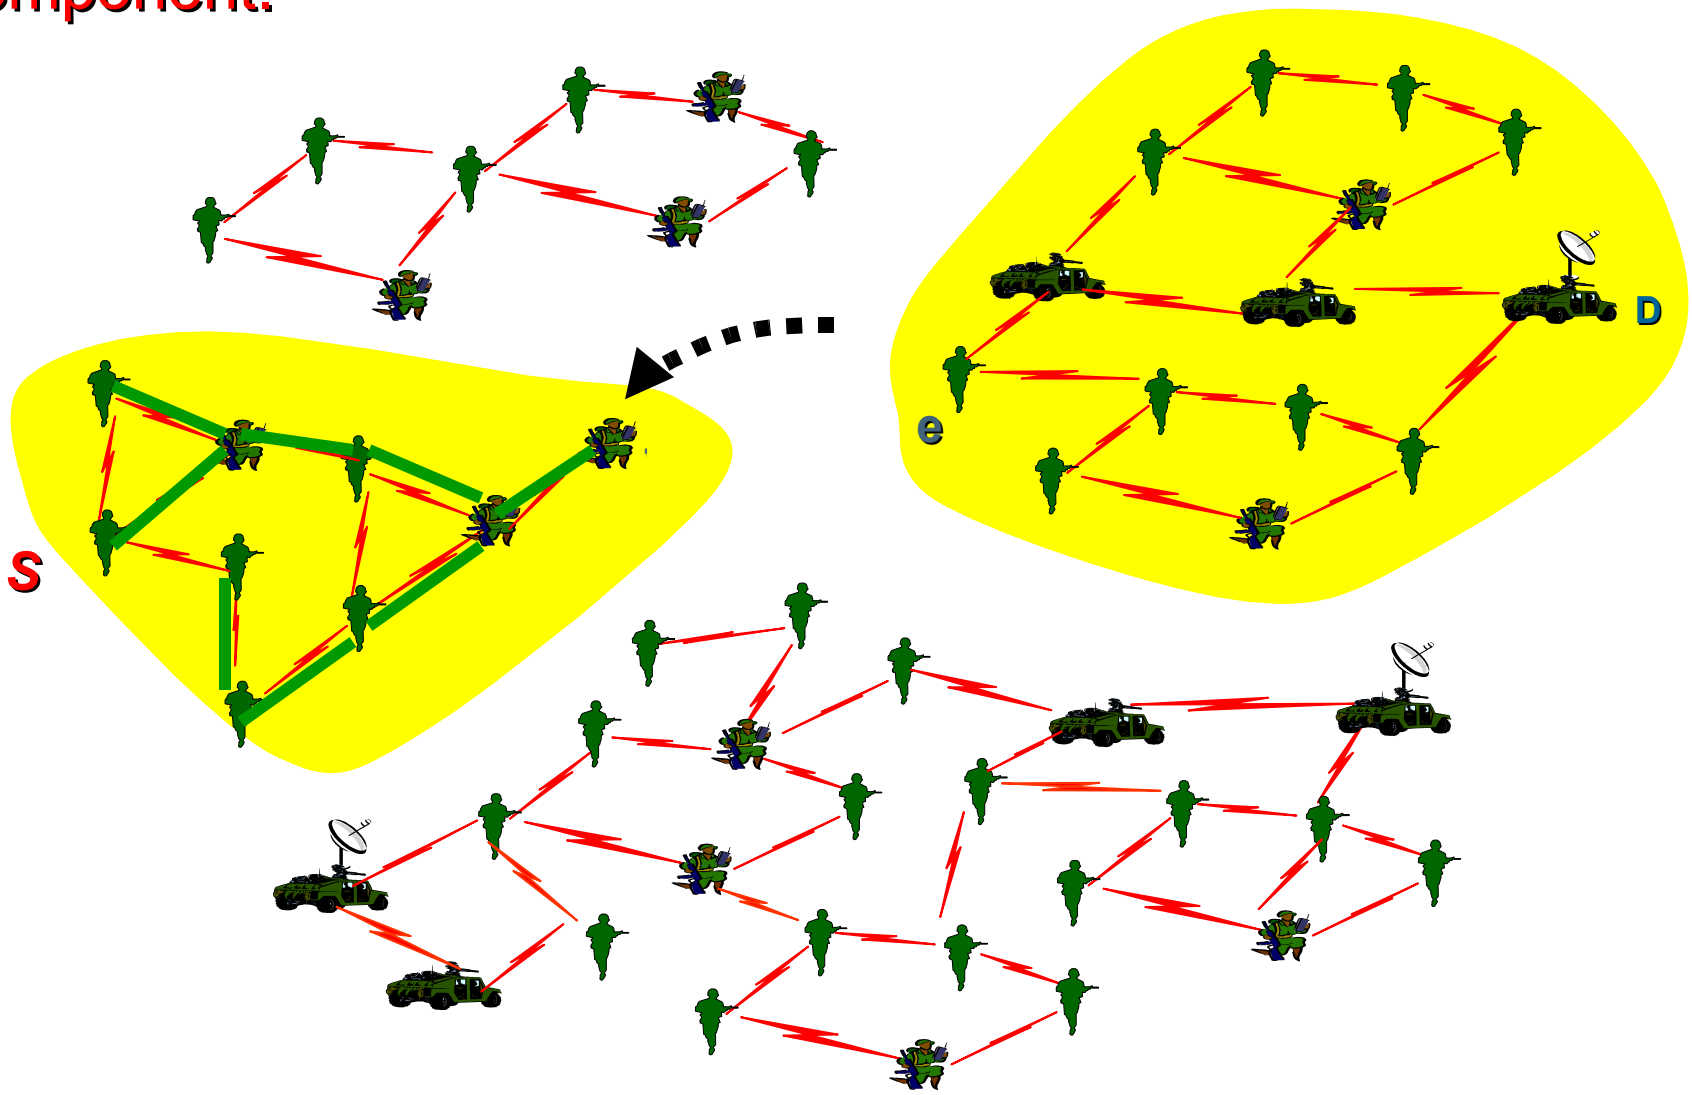

Example

Node S starts sending messages to D through f, which may find a better steward for D or a node with a path to D.

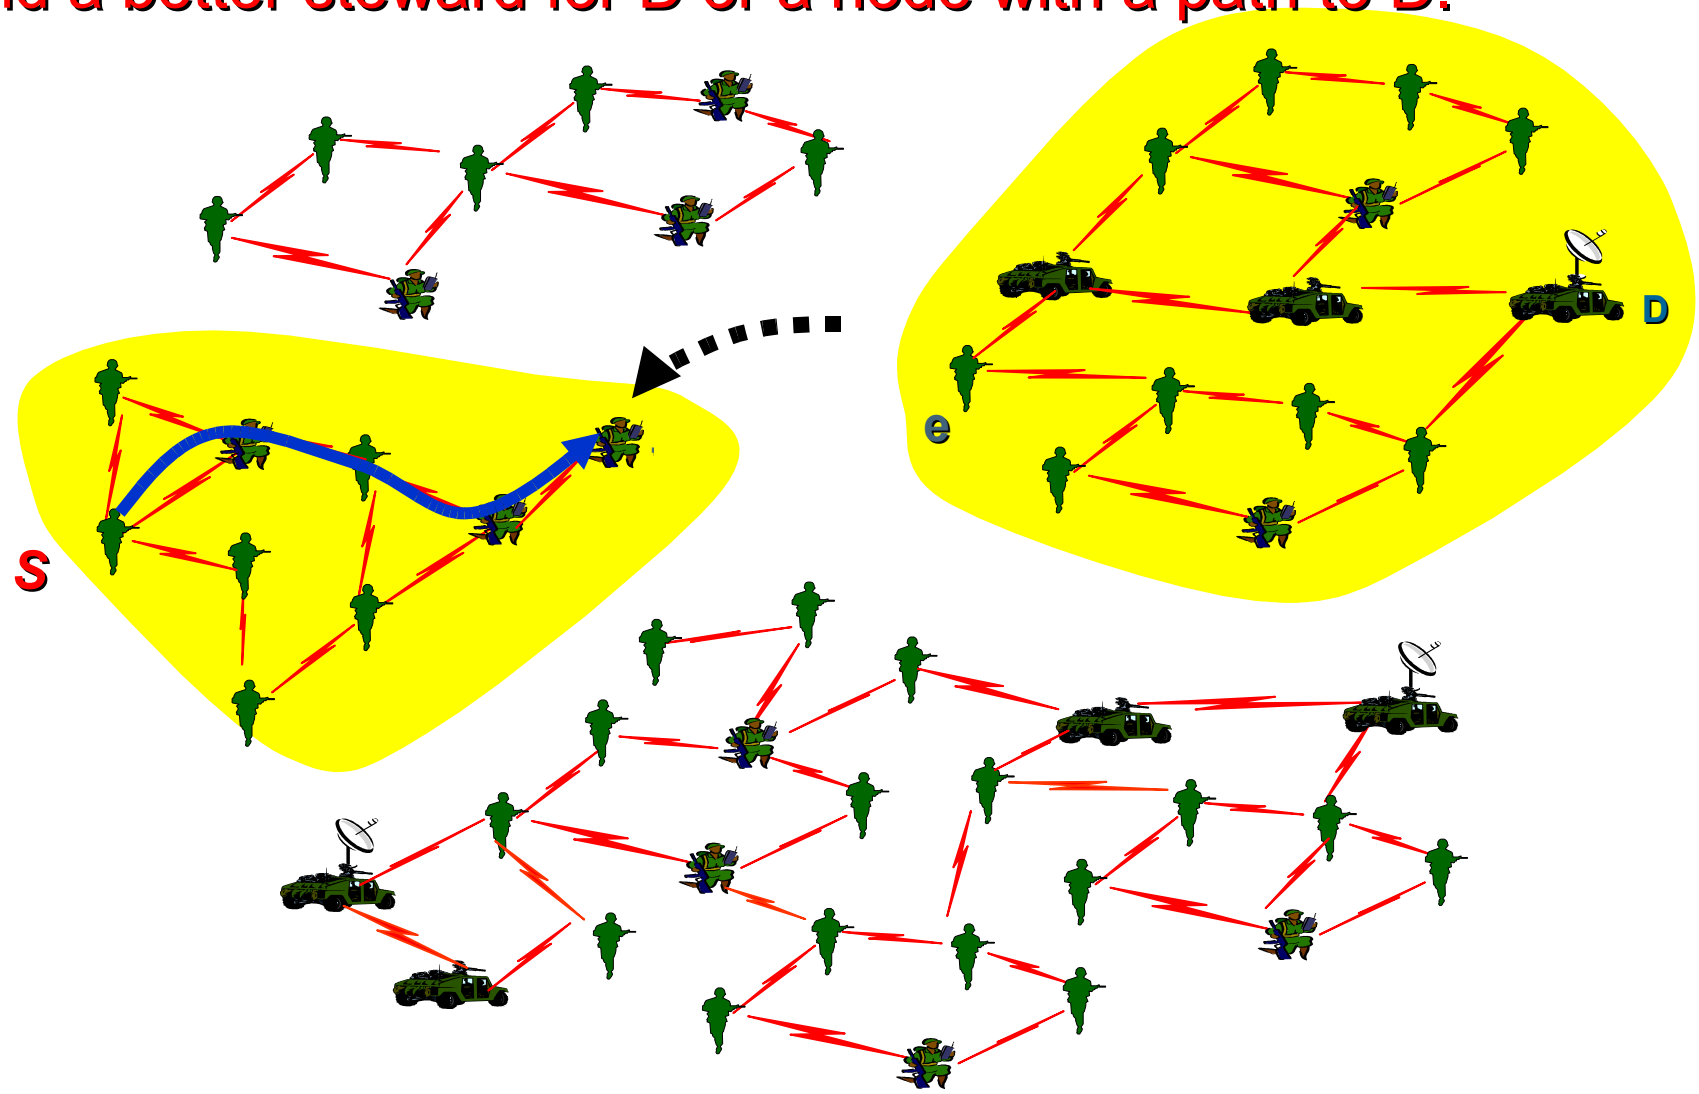

SCIP Impact (Number of routing entries)

- 100 nodes in a (10x10) gridded mobility scenario
- SCIP reduces routing table size proportional to number of sources/sinks and time-based diameter of network.

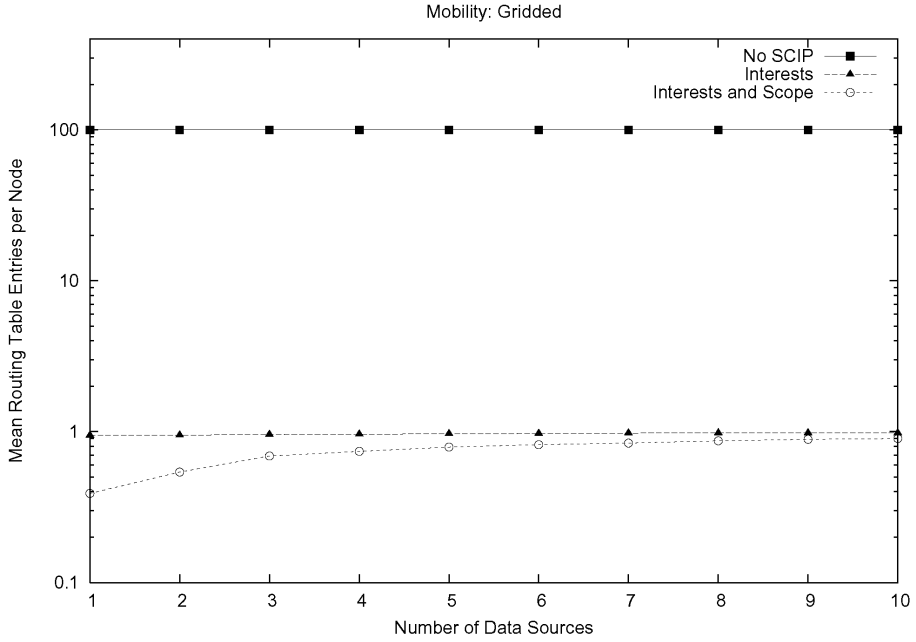

One destination

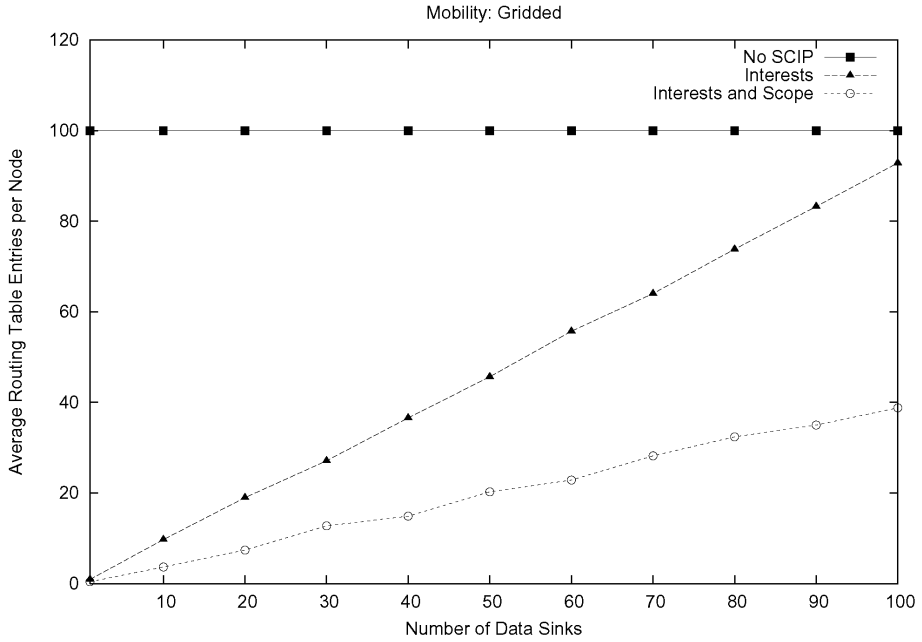

One source per destination

Simulations

- Simulated with two mobility scenarios from real trace data: Dartmouth's CRAWDAD (100 most mobile nodes in October 2004) and UMassDieselNet (30 buses, one day's worth of trace data).
- Provides delivery rates close to that of Epidemic routing, while overhead remains small (independent of buffer size, density, network size).
- Performs best in situations constrained by storage space or bandwidth.

StAR/SCIP Performance

- Dartmouth laptop mobility trace simulation with varied number of laptops, 20 randomly chosen flows.
- Successor forwarding: StAR with message sent to all loop-free successors, not just one

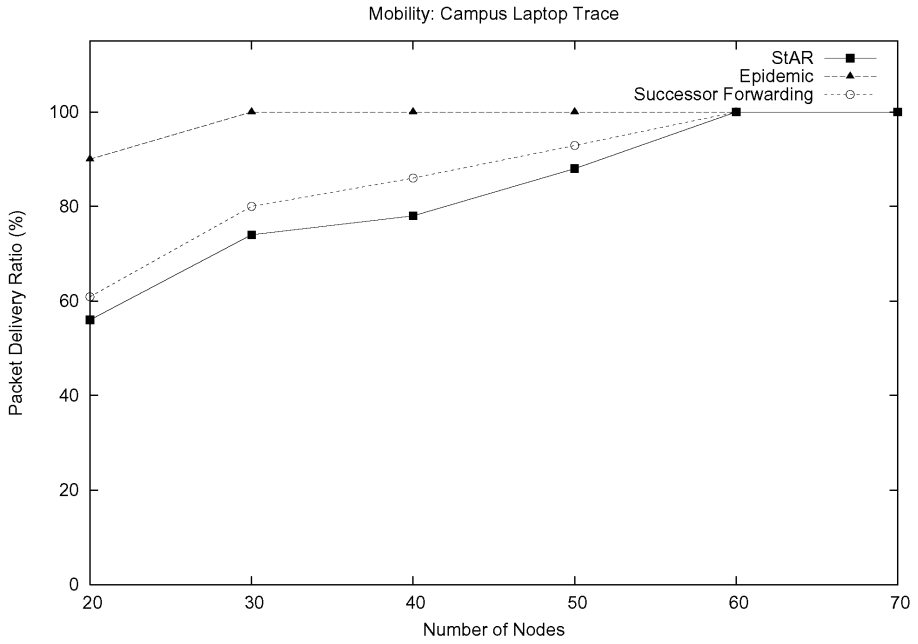

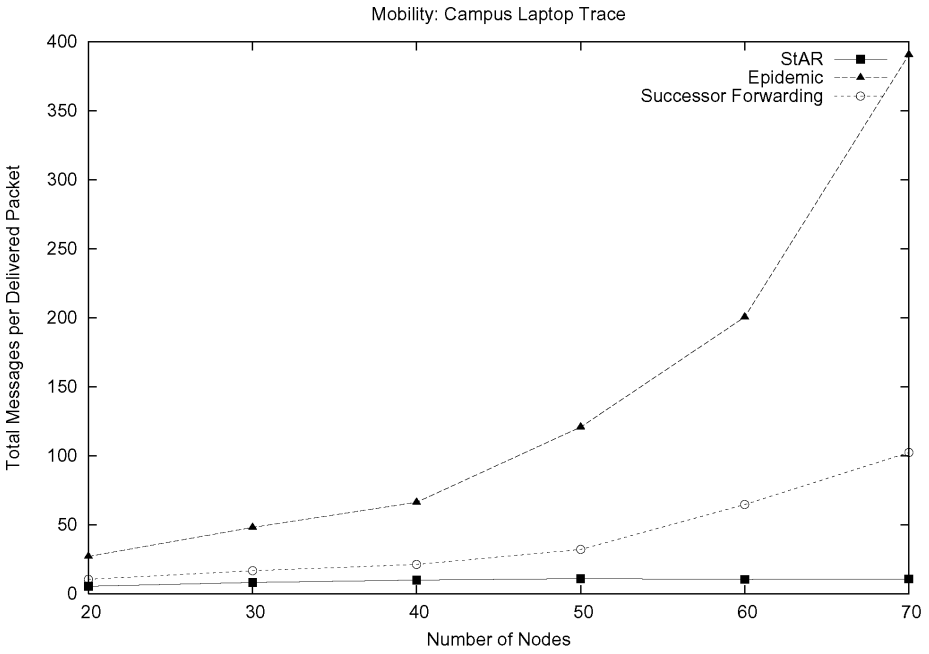

StAR/SCIP Performance

- UMass scheduled bus routes with varied storage space:

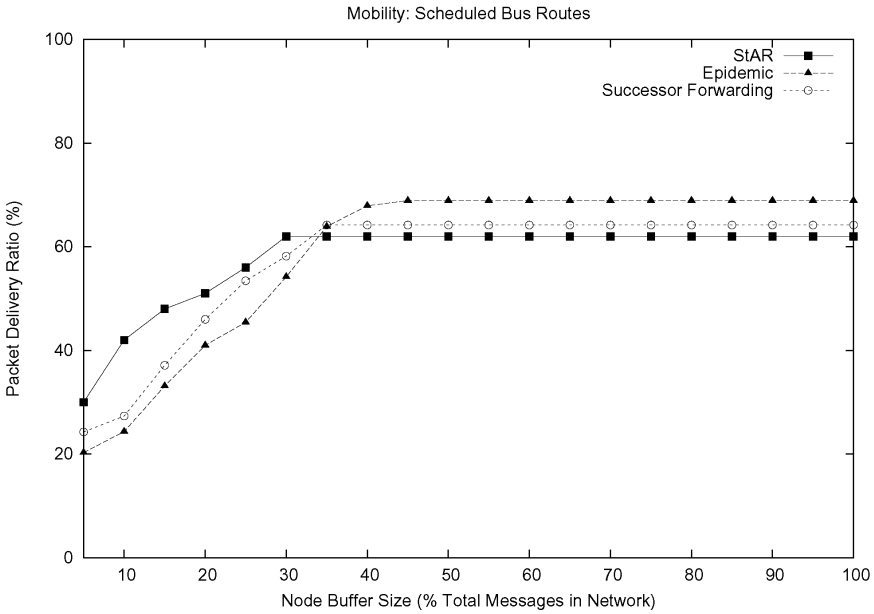

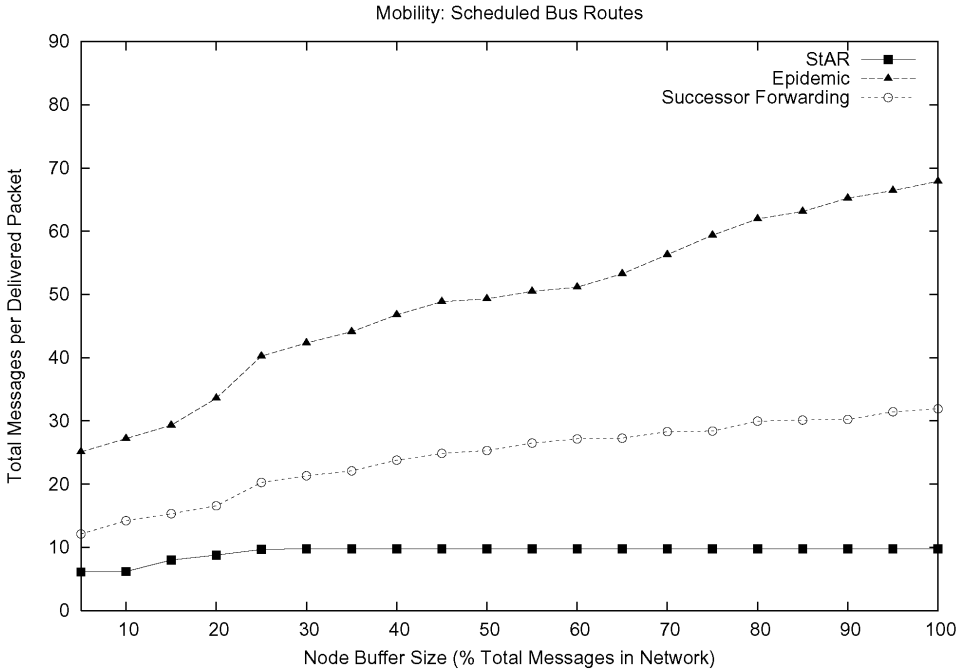

First Next Steps

- Interest has nothing to do with MAC or IP addresses or specific nodes:
  - Use functional and content names
- Use of well known names and stewards as rendezvous points
- Much more efficient schemes to scope the dissemination of *interests* and the existence of destinations are possible!
- Content replication/dissemination with scoping
- Multiple constraints and policies
  - Not all nodes and destinations are equal

The Opportunity: A New Kind of Network

"Store-process-forward" networking;

Process and storage of content inside the network

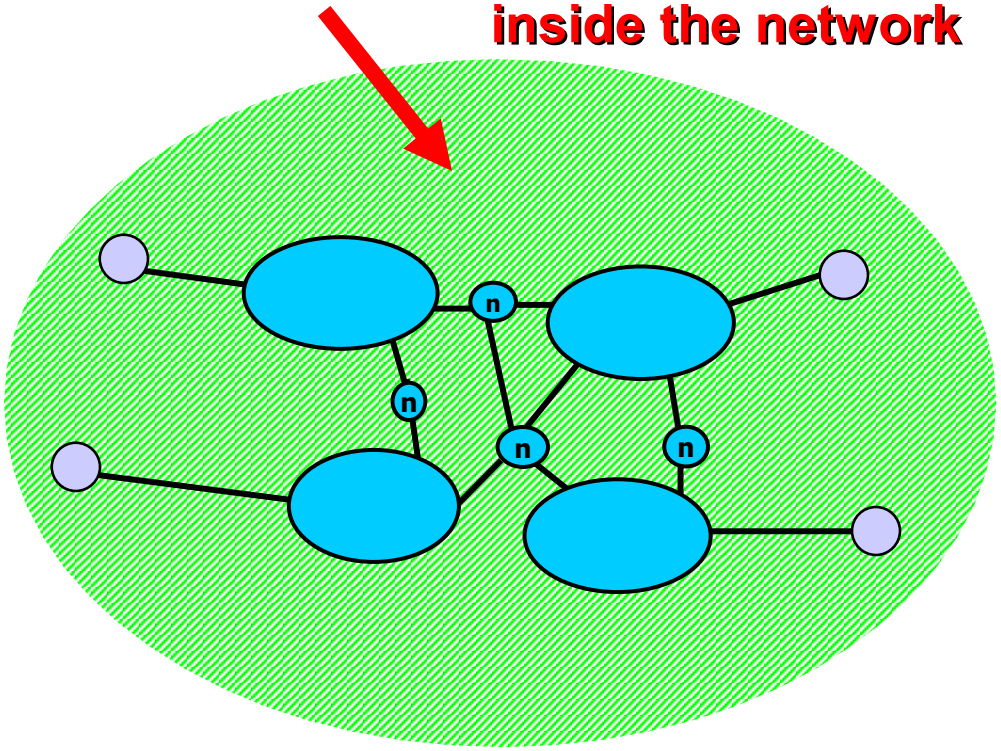

A richer "instruction set" for packet switching that takes advantage of context

New routers store and process content

Names of content, not host addresses, used as the entities for routing

Consumers and providers of content collaborate based on their context

Questions

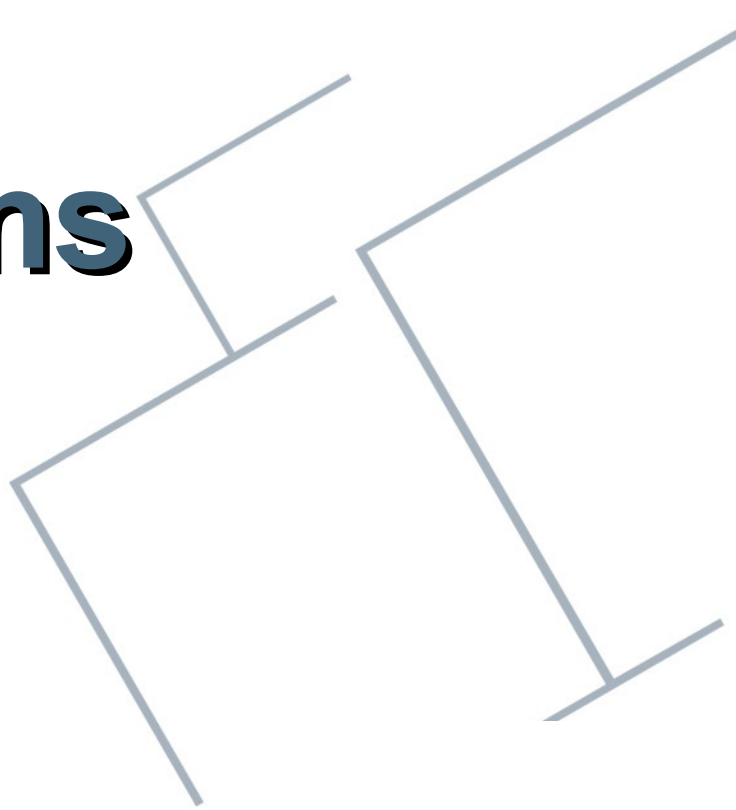

IP Internet Approach

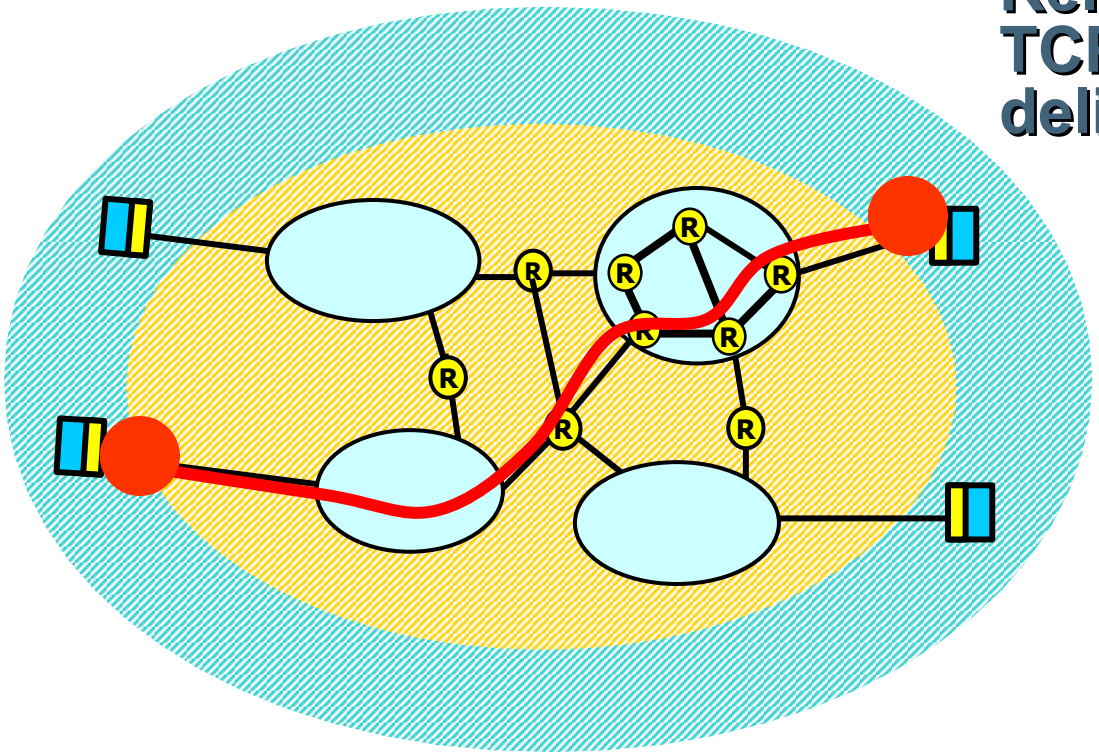

Reliable connections (using TCP) for reliable byte delivery between two hosts

Reliable content delivery via connections between specific hosts is wasteful (>99% use of today's networks is for entities to acquire named chunks of data (like web pages or email messages)

- **Popular sites are hotspots and prone to congestion**
- **Poor reliability from dependence on a channel to the data source**
- **Poor utilization of computing and storage resources in the network**
- **End-to-end connectivity may not be there**

Connection requires connectivity and a bandwidth-delay product that permits feedback.

Flow and congestion control assumes a sender-receiver session against all others.