Age Matters :
Efficient Route Discovery Using Encounter Ages

Henri Dubois-Ferrière, Matthias Grossglauser, Martin Vetterli
School of Computer and Communications Sciences
Swiss Federal Institute of Technology (EPFL) Lausanne

Mobihoc 2003, Annapolis, MD
June 3rd 2003

Context

FRESH (FResher Encounter SearcH) is an algorithm for *route discovery* in ad hoc networks.

Applicable to protocols which work :

- On demand
- Blindly (no location info, no GPS)

Context

FRESH (FResher Encounter SearcH) is an algorithm for *route discovery* in ad hoc networks.

Applicable to protocols which work :

- On demand
- Blindly (no location info, no GPS)

Most relevant to : *Large* networks (several hundred nodes at least) of *mobile* nodes.

Context

FRESH (FResher Encounter SearcH) is an algorithm for *route discovery* in ad hoc networks.

Applicable to protocols which work :

- On demand
- Blindly (no location info, no GPS)

Most relevant to : *Large* networks (several hundred nodes at least) of *mobile* nodes.

Design Goals :

- Increase scalability of current routing protocols.
- General algorithm that can be integrated in existing protocols :
  - Not tied to one particular protocol or approach.
  - Can leverage existing work - avoid doing YAP (Yet Another Protocol).

Age Gradient : Make Mobility work for you !

- Mobility Increases Capacity [Grossglauser, Tse Infocom 2001]
- EASE [Grossglauser, Vetterli Infocom 2003] : Scalable location service. Nodes remember *where* and *when* they met other nodes.

Age Gradient : Make Mobility work for you !

- Mobility Increases Capacity [Grossglauser, Tse Infocom 2001]
- EASE [Grossglauser, Vetterli Infocom 2003] : Scalable location service. Nodes remember *where* and *when* they met other nodes.

FRESH shows that routing can gain in efficiency even if nodes only remember *when* they met others. This makes use of the *Age Gradient* :

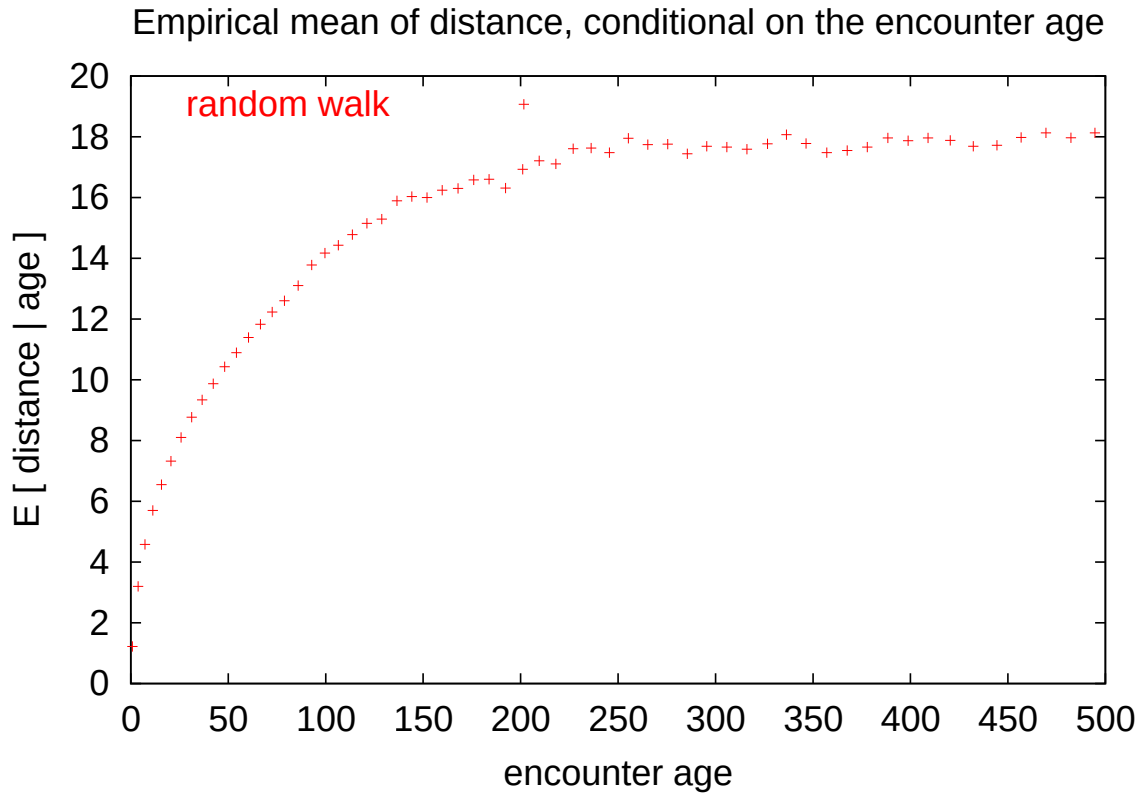

Motivation : Scaling Issues

Scaling to larger networks requires to reduce the routing overhead.

- Route discovery inevitably resorts to some form of flooding (global, probabilistic, expanding ring search, etc).
- Flooding is expensive, e.g. :
  Up to 90% of routing overhead in AODV [Das, Perkins, Belding-Royer, Infocom00]

Motivation : Scaling Issues

Scaling to larger networks requires to reduce the routing overhead.

- Route discovery inevitably resorts to some form of flooding (global, probabilistic, expanding ring search, etc).
- Flooding is expensive, e.g. :
  Up to 90% of routing overhead in AODV [Das, Perkins, Belding-Royer, Infocom00]

"The fat is in the flood" ... so how do we trim it down ? By restricting the flood to a smaller area and to fewer nodes !

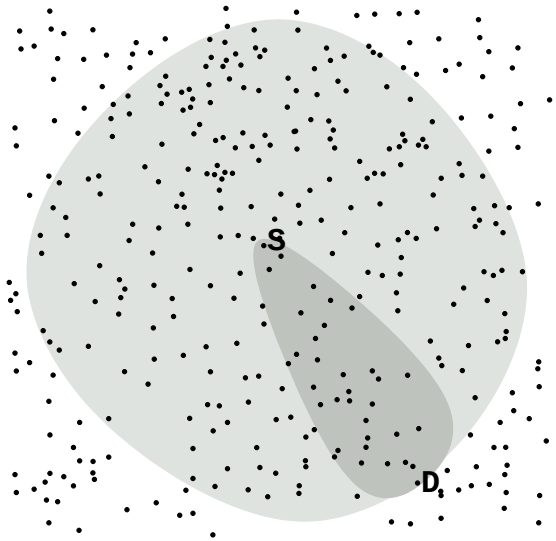

Encounters

Definition : An *Encounter* between two nodes A and B happens when these nodes are one-hop neighbors. (Nodes are one-hop neighbors when within xmit range, without relaying.)

Each node keeps a table of its *last encounter times* with all other nodes.

Encounter table at node S : T_S.

1. $T_S(D)$: time elapsed since S, D were last one-hop neighbors
2. $T_S(D) = \infty$ if S, D were never one-hop neighbors
3. $T_S(S) = 0$

How do nodes "know" when they make encounters ?

1. Overhearing any packet in promiscuous mode
2. "Hello" broadcast messages (Many schemes possible : periodic, adaptive, ..)

Baseline Intuition

1. *Several local searches are cheaper than one global search.*

Therefore, instead of searching for the destination :

Search for any node that is closer to the destination, and repeat from there.

Baseline Intuition

1. *Several local searches are cheaper than one global search.*

Therefore, instead of searching for the destination :

Search for any node that is closer to the destination, and repeat from there.

2. *How to find a node which is closer to the destination ?*

Search for a node which was a neighbor of the destination more recently than myself.

FRESH Algorithm

Source S establishes route to destination D :

Start at node S.
S floods until reaches a node $A1$ s.t $T_{A1}(D) < T_S(D)$
$A1$ floods until reaches a node $A2$ s.t $T_{A2}(D) < T_{A1}(D)$
$\ldots$
until D is reached (for which $T_D(D) = 0$)

Notes :

1. "Flood until" $\equiv$ some form of expanding ring search.
2. Encounter age decreases at each iteration. Therefore FRESH terminates and is loop-free.
3. Each iteration is equivalent to one route discovery with a standard protocol (AODV, DSR).
4. FRESH works equally well with distance-vector routing as with source-routing.

Example FRESH Route

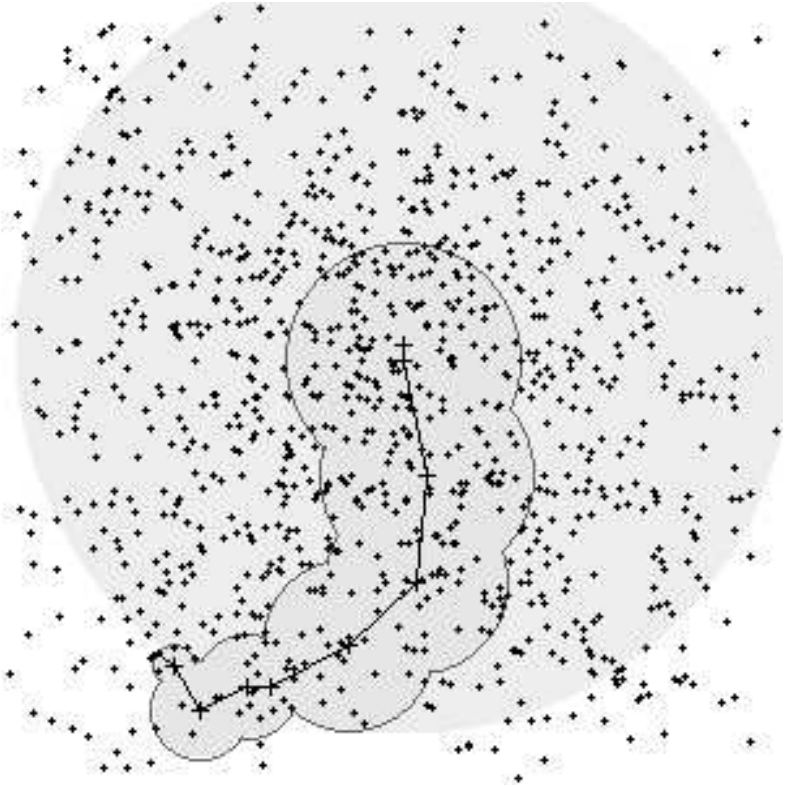

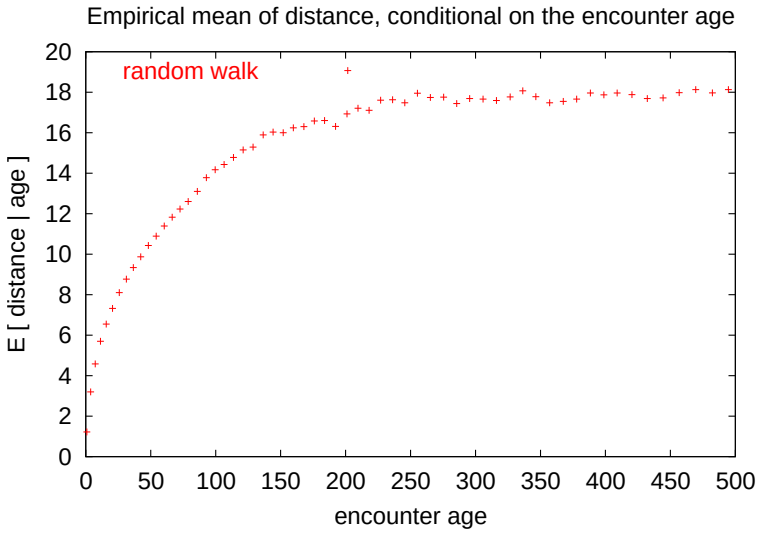

Source in center, Dest at bottom.

Dark surface : area covered by union of (minimal) fresher encounter searches.

Light surface : area covered by (minimal) direct route discovery search (existing protocols).

Performance Metrics

Route Quality : Ratio of effective route length to optimal route length

Search Cost : Total area flooded (Note : proportional to # pkt xmits, # nodes involved, etc)

Remarks :

1. These two metrics do not go hand in hand !
2. FRESH trades off some route quality in exchange for lower search cost.
3. Route quality vs. search cost weighting depends on session duration.

Simulation Details

Simulation Details :

- 64000 nodes
- Square surface, sized to have unit density
- Random walk and waypoint models
- Warmup runs to 40% encounter ratio
- Ideal flooding approximation
- Routes computed sequentially - no cross-traffic.
- Simple MAC Layer

Simulation Details

Simulation Details :

- 64000 nodes
- Square surface, sized to have unit density
- Random walk and waypoint models
- Warmup runs to 40% encounter ratio
- Ideal flooding approximation
- Routes computed sequentially - no cross-traffic.
- Simple MAC Layer

`mws` (multihop wireless simulator) :

- Design similar to `ns-2/Monarch`.
- Lean and highly efficient, optimized for large-scale simulations.
- Public release expected Fall 2003 ... send email if interested !

  `(henri.dubois-ferriere@epfl.ch)`

Simulation Results : Search Cost

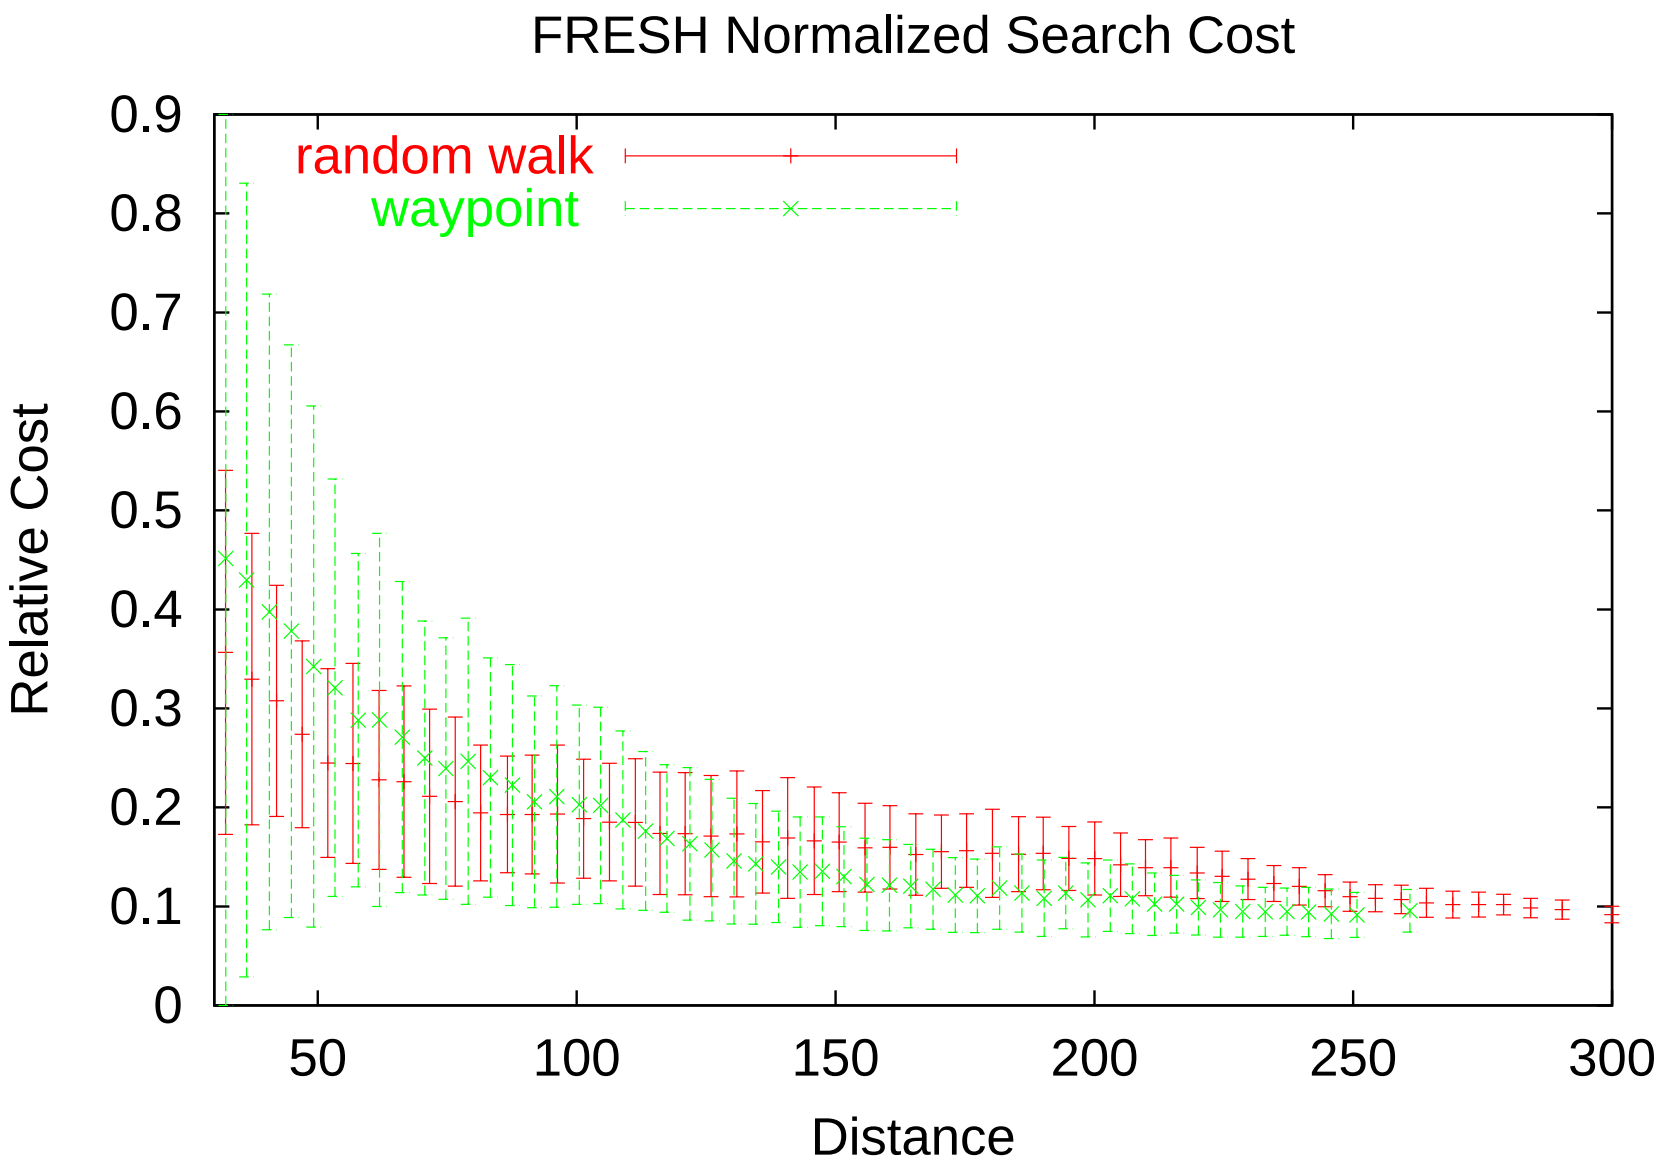

Simulation Results : Route Quality

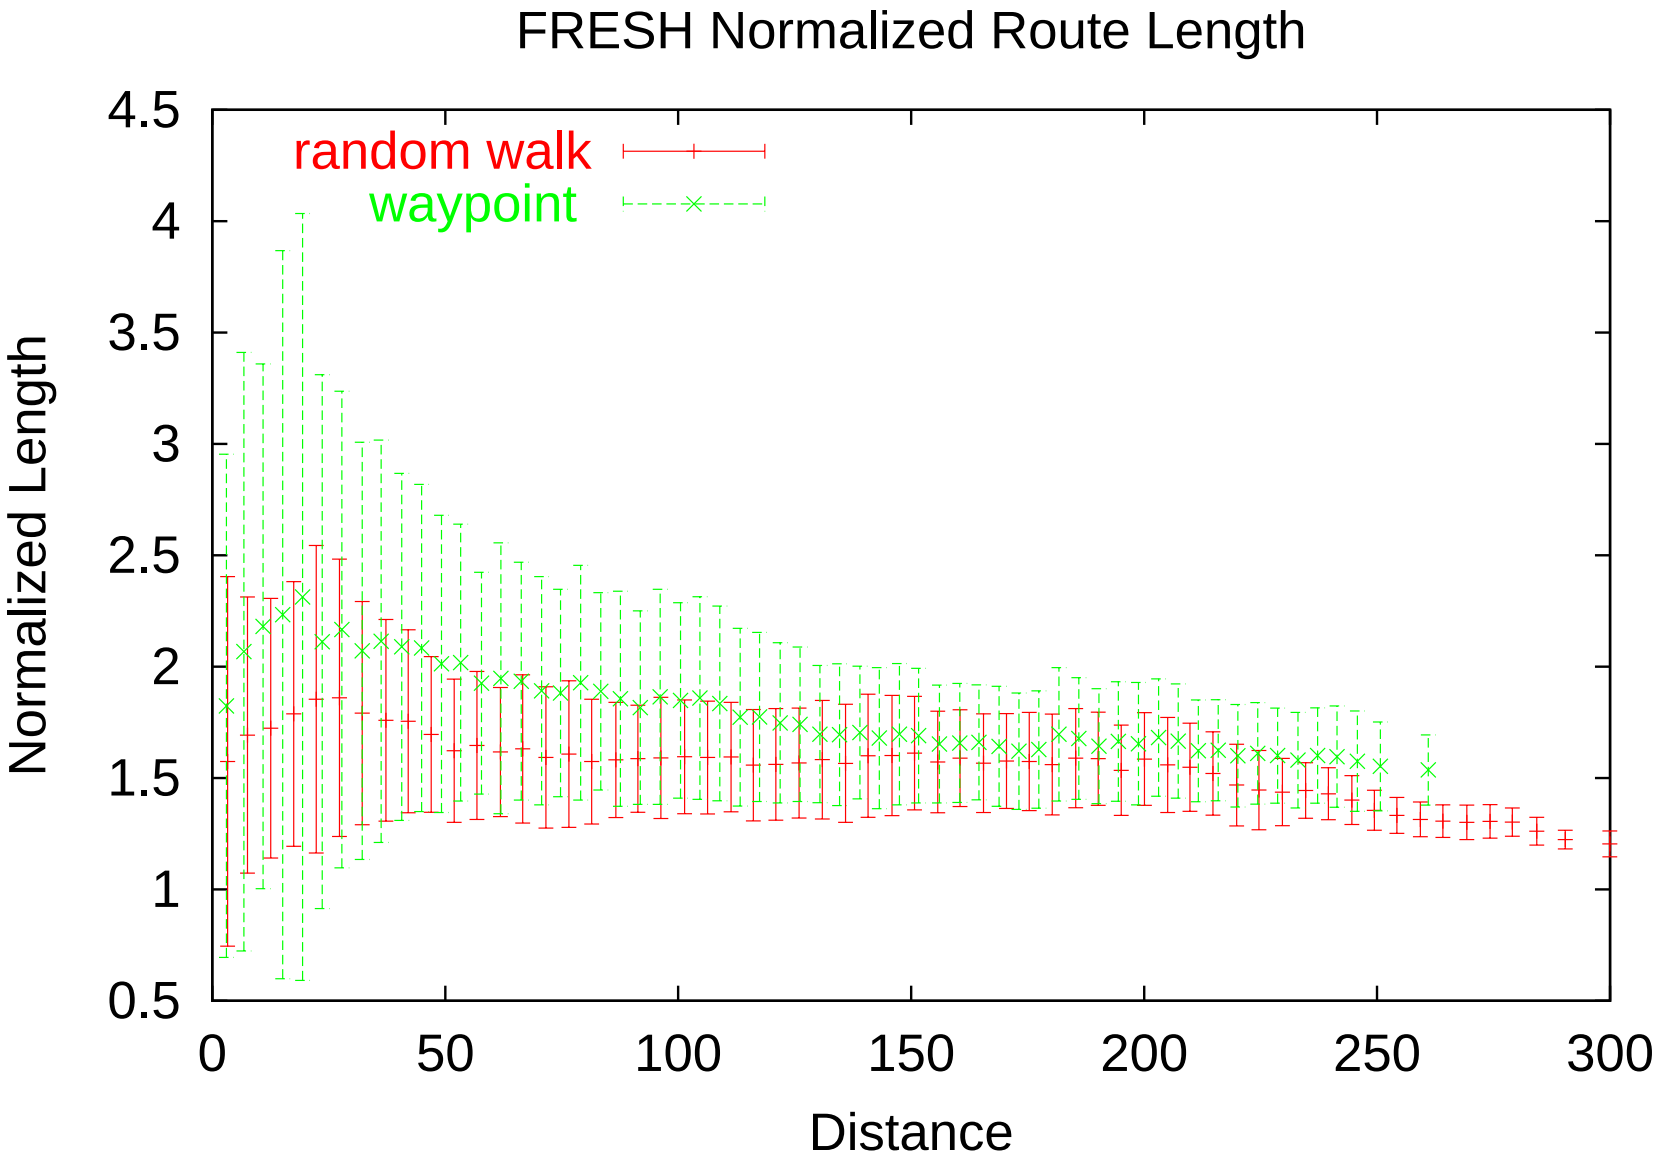

Space vs. Time Routing

1. "Traditional" protocols *route through space* :
   Use a *spatial* distance metric (e.g., hops) to advance towards destination.
2. FRESH *routes through time* :
   Use a *temporal* distance metric (last encounter age).

Are both approaches exclusive ? No : Space-Time Routing

- Jointly uses spatial and temporal metrics.
- Parametric algorithm - interpolates anywhere between purely on-demand to purely proactive
- Generalizes FRESH, on-demand distance-vector routing (similar to AODV), and proactive distance-vector routing (similar to DSDV).

Initial results promising.. more details soon !

FRESH - Conclusions

1. Efficient
   - Tenfold gain despite conservative cost metric.
   - Some obvious optimizations will further improve efficiency.
2. General
   - Can be integrated to current protocols (AODV, DSR,..) - leverage existing work !
   - Loop-free
3. Exploits node mobility
   - See EASE [Grossglauser, Vetterli Infocom 2003]
4. Trade off route quality for reduced search cost.
   - Where is the optimal tradeoff ?
   - Can we adjust the tradeoff point based on session duration ?