

olsr.org

'Optimized Link-State Routing' and beyond

December 28th, 2005

Elektra [www.scii.nl/~elektra](http://www.scii.nl/~elektra)

# 'Optimized Link-State Routing' and beyond


# December 28th, 2005

Elektra [www.scii.nl/~elektra](http://www.scii.nl/~elektra)

# Introduction


- Olsr.org is aiming to an efficient open-source routing solution for wireless networks
- Work is currently based on the Olsr-protocol suggested by RFC3626
- There is not much left from RFC3626 now, though. You'll see why...

# Idea: Multipoint to Multipoint Networking



- A.k.a. Mesh-Networking
- Wireless network based on 802.11 nodes, operating in Ad-Hoc-Mode
- Cover large areas: A and D talk via B and C


# 802.11 Managed Mode doesn't allow this\*



- A talks to B via central Access Point
- C cannot talk to B or A – although B would be in range of C's Wifi Link
- D and B have to use the AP as relay, thus speed is only 50%


\*) when operating with a single interface

# 802.11 Ad-Hoc Mode



- Nodes talk directly
- Decentralized & scalable when routing is applied

# Proactive Link-State Routing



- Link Detection via Hello-Broadcasts
- A and B notice each other

- Topology Information flooding
- A says 'I see B', B says 'I see A, C, D' a.s.o.

# Flooding of Topology Information


The diagram illustrates the flooding of topology information in a network. It shows four scenarios of node interactions (A, B, C, D) across different network topologies. Nodes are represented by colored circles: blue for A, green for B, yellow for C, and red for D. Arrows indicate the direction of information flow. The background features a complex network graph with various IP addresses and link costs.

olsr.org – 'optimized link state routing' and beyond

7




# Topology Message Flooding



- All neighbors retransmit messages all over the network
- Bandwidth usage
- Wasting CPU-Cycles
- Collisions



# Dijkstra's Algorithm




	A	B	C	D
A	*	*	B	B
B	*	*	*	*
C	B	*	*	*
D	B	*	*	*

- Everybody knows everybody else and their links
- Routing table: Dijkstra's Algorithm for shortest paths

# OLSR Basics

- INRIA-Draft specified by RFC 3626
- Proactive, using Dijkstra's Algorithm
- Communication via UDP broadcasts
- Multiple OLSR messages per UDP packet
- Validity time in OLSR messages
- Information discarded by timeouts
- Introduced new ideas that were meant to reduce protocol overhead and increase stability: Hysteresis, MultiPointRelays

# RFC3626 Idea: Reducing Overhead



- Only selected neighbors (Multi-Point Relays, MPRs) retransmit messages
- Select MPRs such that they cover all 2-hop neighbors
- 2-hop neighbors taken from neighbors' HELLO messages
- Does not work in real-life! Reduces redundancy and stability!

# Issues in the INRIA-Draft

- Adds new and unnecessary message class of MPRs
- Still optimizes for lowest Hop-Count
- Discards links to neighbours by Hysteresis
- Reduces topology information redundancy
- Every node floods the whole network (at least all MPRs)
- Breaks the KISS-Attitude!

# Real-life results of RFC 3626

- Routing table breaks down all the time
- Prefers routes with shortest path, low bandwidth and no stability
- Routing loops occur very often

# Lessons learned by using RFC3626

- A mesh is a boiling kettle with interference and collisions
- Theoretical solutions in simulations are unlikely to work in real life scenarios
- Make it work. Make it stable. Worry later about optimizations routing the whole universe in one subnet...
- Proactive routing algorithms depend on synchronized information. Transmissions must be redundant.
- New message types introduce new headaches.

# What we did...

- Disable Hysteresis in the configuration file
- Disable MultiPointRelay selection
- Implement route calculation depending on packet loss (LQ-ETX)
- Implement fish-eye mechanism for forwarding of topology information (Link-Quality-Fish-Eye. New in olsr-0.4.10)

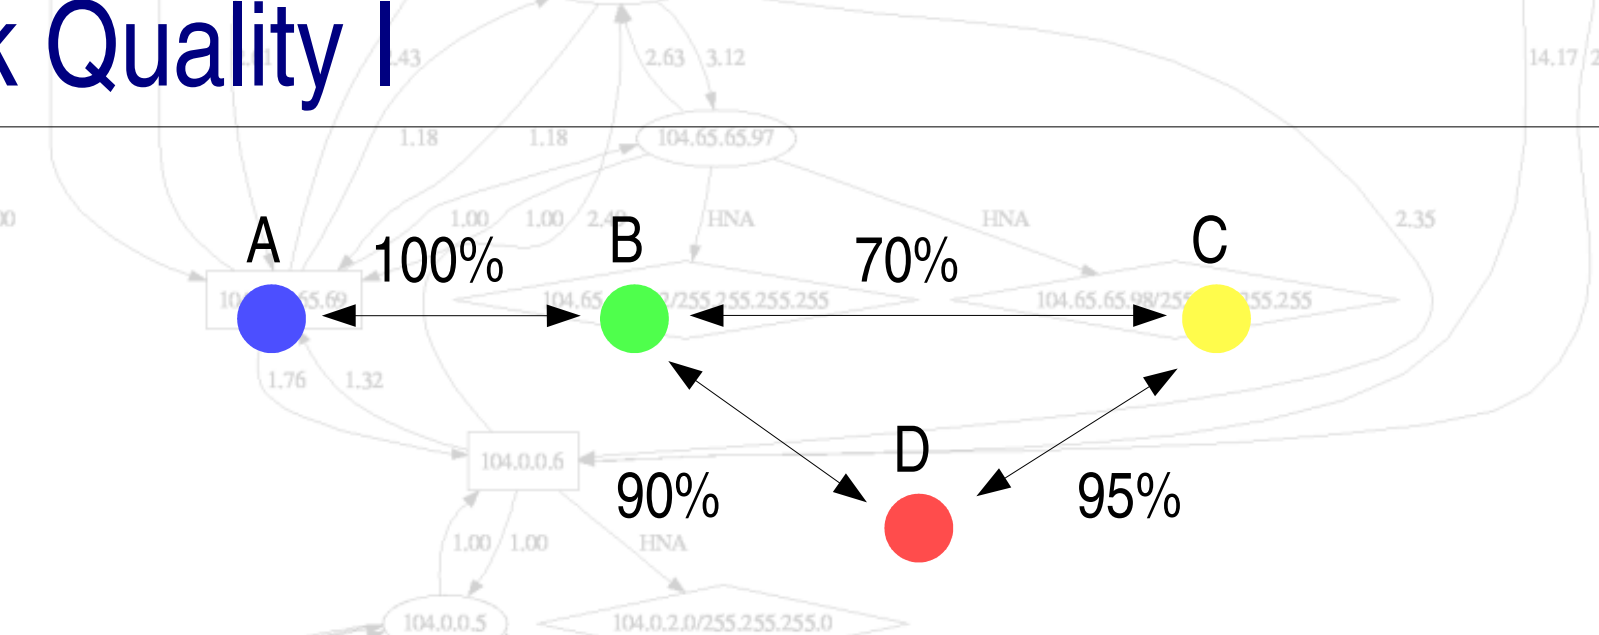


# Link Quality I


- OLSR minimizes hop count, hence favors longer (lossier) links
- Alternative – minimize packet loss
  - A – B – C with 70% path quality
  - A – B – D – C with 85% path quality
- Other metrics – latency, throughput, ...

olsr.org – 'optimized link state routing' and beyond

16

- # Link Quality I
- 
- OLSR minimizes hop count, hence favors longer (lossier) links
  - Alternative – minimize packet loss
    - A – B – C with 70% path quality
    - A – B – D – C with 85% path quality
  - Other metrics – latency, throughput, ...
- olsr.org – 'optimized link state routing' and beyond
- 16

# Link Quality II

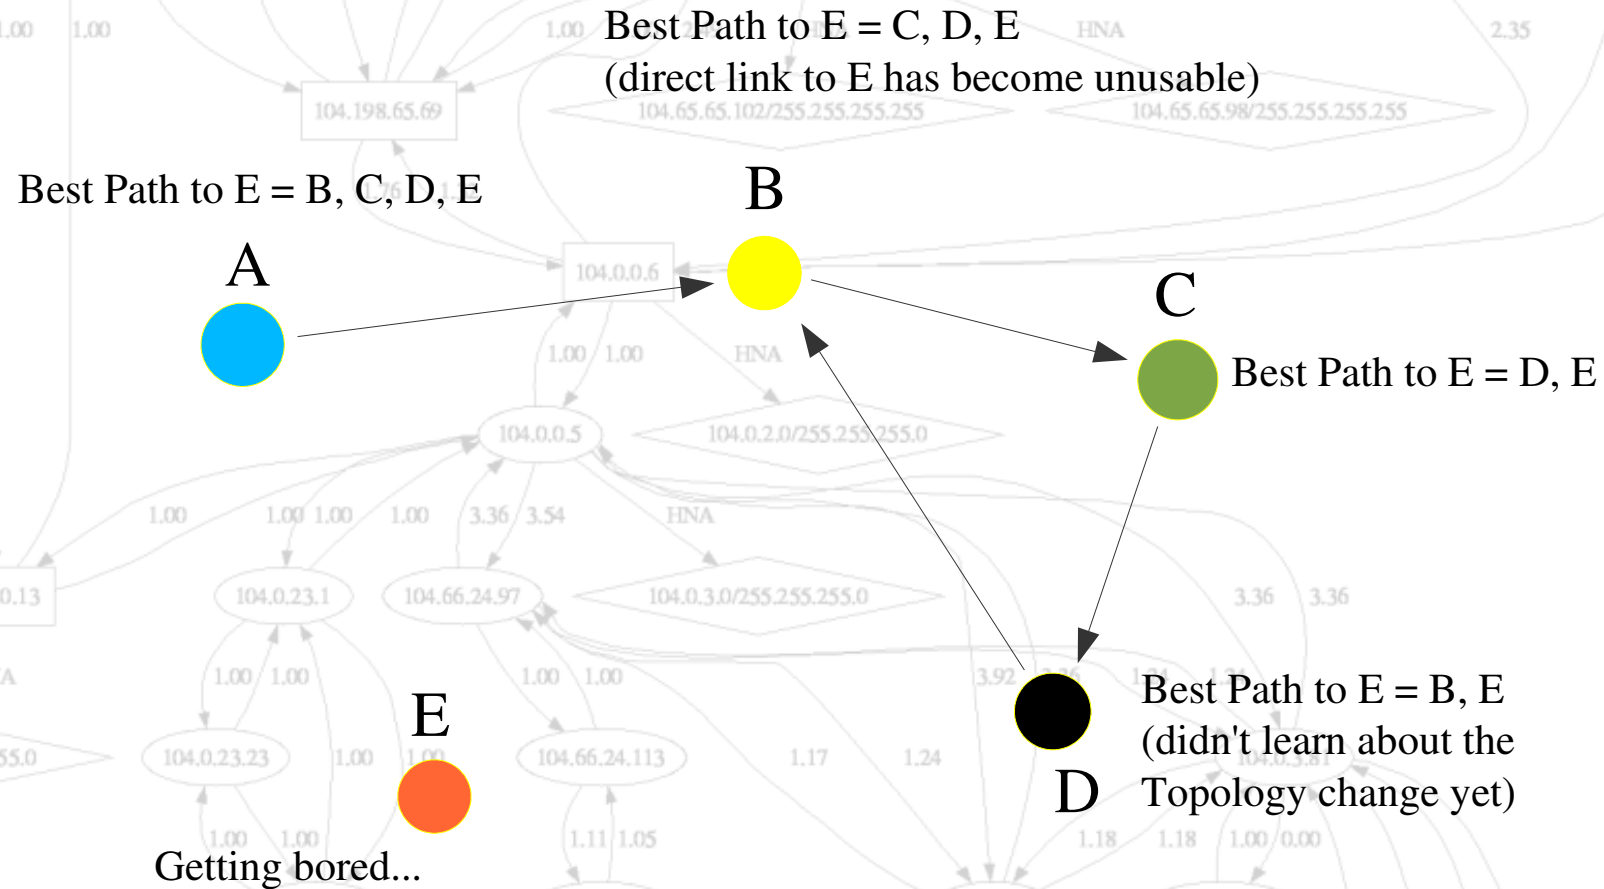


- Minimize Expected Transmission Count (ETX)
- Retransmission – packet or acknowledgment lost
- Packet loss among recent x HELLO messages
- $LQ_1 = 90\%$ ,  $LQ_2 = 70\%$
- $ETX = 1 / (LQ_1 \times LQ_2) = 1 / 0.63 = 1.59$

# Result: Olsr.org works

- Many people successfully share DSL-Lines with their mesh.
- Networks with up to 150 nodes work well
- Still issues under high traffic load – as links saturate routing loops occur.
- Networks that don't saturate their Wifi-Links are not affected.
- The Berlin mesh with more than 250 routes pushes small CPUs to the limit


# A typical routing loop



# Addressing the routing-loop issue

- Occurs when topology information is not in sync
- Loops happen amongst adjacent nodes
- Interference causes topology information loss
- Payload traffic causes interference
- Topology information must be redundant and sent often, more often than Hello-messages to provide information timely
- MultiPointRelays don't help

# Link Quality Fish Eye



- Broadcast topology messages with small TTL often
- Send messages with large TTL seldom
- Distant nodes have hazy view – sufficient
- Saving CPU-Cycles
- Saving Collisions



# Implementation

- olsrd 0.4.10 – [www.olsr.org](http://www.olsr.org)
- Linux, \*BSD, Mac OS X, Windows
- Reasonably stable – Berlin and Amsterdam  
(More than 200 Nodes in Berlin)
- Plug-in interface (OLSR Flooding)
- Web-based monitoring
- Link Quality Fish Eye Algorithm



Thanks for  
your  
attention.

Questions?