B.A.T.M.A.N. V – STATUS UPDATE 07/04/2013

List of currently implemented componentss:

ELP:

- periodic broadcast sending *
- periodic unicast sending (for throughput sampling). One unicast ELP message is sent to each of the neighbours to let the RC algorithm make some throughput computations.
- new ELP neighbour struct definition
- neighbour purging on timeout *
- neighbour dead detection based on OUTDATED_MAX ELPs, but no action is taken (current value is 4 as suggested in the wiki)
- throughput reading via cfg80211 (patch not yet upstream)
    - throughput is the maximum value estimated by the RC algorithm
- default throughput value customisable via sysfs (value used for non wireless interfaces)
- forced per interface throughput value via sysfs

OGMv2:

- re-organization of orig_node/neigh_node structures for multiple
- periodic broadcast sending
- routing logic and OGM processing as described on the OGMv2 wikipage
- metric decremented as follows before forwarding:
    - new_metric = metric / 2 (if incoming iface is wifi)
    - new_metric = metric * (1 - hop_penalty) otherwise
- usage of throughput as metric (it is read from the new ELP neighbour struct)
- customised originator table output
- introduction of get_metric() API for "external" metric usage (GW, IFALT, etc..)
- neigh_node invalidation after OGM_SEQ_RANGE missing OGMs (current value is 5, as suggested in the wiki)
- TT component slightly fixed to become algorithm agnostic
- introduction of orig_node/neigh_node routing API for private operations on: object freeing, add new hard_iface, del hard_iface

Simulated scenarios:

- topologies reported on the "Routing Scenarios" wikipage have been tested with NO packet loss to prove the effectiveness of what has been implemented so far
- In case of link breakage, the time needed to recover and switch to a new path is OGM_SEQ_RANGE * originator_interval (time needed to invalidate the old router).

* = components I found implemented already

Algorithm "cost" comparison (B.A.T.M.A.N. IV vs. B.A.T.M.A.N. V)

Costs are computed on a simple topology made up by N nodes, where each of them can hear each other (so each node has N-1 direct neighbours).

- Variables description
  - N number of nodes in the network
  - O_{bat4} B.A.T.M.A.N. IV originator interval
  - O_{bat5} B.A.T.M.A.N. V originator interval
  - E_{uni} broadcast ELP interval (for neighbour discovery)
  - E_{brd} unicast ELP interval (for throughput sampling)
- Packet rate model (restricted to the aforementioned topology)

We use the following models to "count" the number of packets generated by each algorithm and later we determine which parameter we can tune in order to make B.A.T.M.A.N. V behave better than B.A.T.M.A.N. IV.

- B.A.T.M.A.N. IV:

pkt/s = $N\left(\frac{1}{O_{bat4}}+\frac{N-1}{O_{bat4}}\right)$ each node sends 1 and rebroadcast N-1 OGMs

$$N\left(\frac{N}{O_{bat4}}\right)$$

$$\frac{N^2}{O_{bat4}}$$

- B.A.T.M.A.N. V:

pkt/s = $N\left(\frac{N}{O_{bat5}}+\frac{1}{E_{brd}}+\frac{N-1}{E_{uni}}\right)$ each node sends N OGMs (own + rebroadcast), 1 broadcast ELP and N-1 unicast ELP

$$\frac{N^2}{O_{bat5}}+\frac{N}{E_{brd}}+\frac{N^2-N}{E_{uni}}$$

Assuming a constant value for the following variables:

- N 5 nodes
- O_{bat4} 1 second
- E_{brd} 0.5 seconds

we can simplify the above formulas to the following:

- B.A.T.M.A.N. IV

pkt/s = $\frac{25}{1}=25$

- B.A.T.M.A.N. V

pkt/s = $\frac{25}{O_{bat5}}+\frac{5}{0.5}+\frac{25-5}{E_{uni}}=\frac{25}{O_{bat5}}+10+\frac{20}{E_{uni}}$

As we can see, the packet rate for B.A.T.M.A.N. IV is now constant and this gives us the chance to play with the values of O_{bat5} and E_{uni} and let B.A.T.M.A.N. V generate less overhead than its predecessor.

As we can see from the results, the unicast ELP interval plays an important role as much as the originator interval and this can be explained by the fact that both are driven by a function in the square of the number of nodes in the network.

The first (and easiest) attempt we can set up consist in letting the unicast ELP interval vary among a set of values and observe the behaviour of the number of messages in function of O_{bat5}.
So we let E_{uni} vary in the following set *E = { 0s (disabled), 0.5s, 1s, 5s }*.
<u>To simplify the reading of the following results, values have been computed **per minute**.</u>

- **0s (Disabled)**

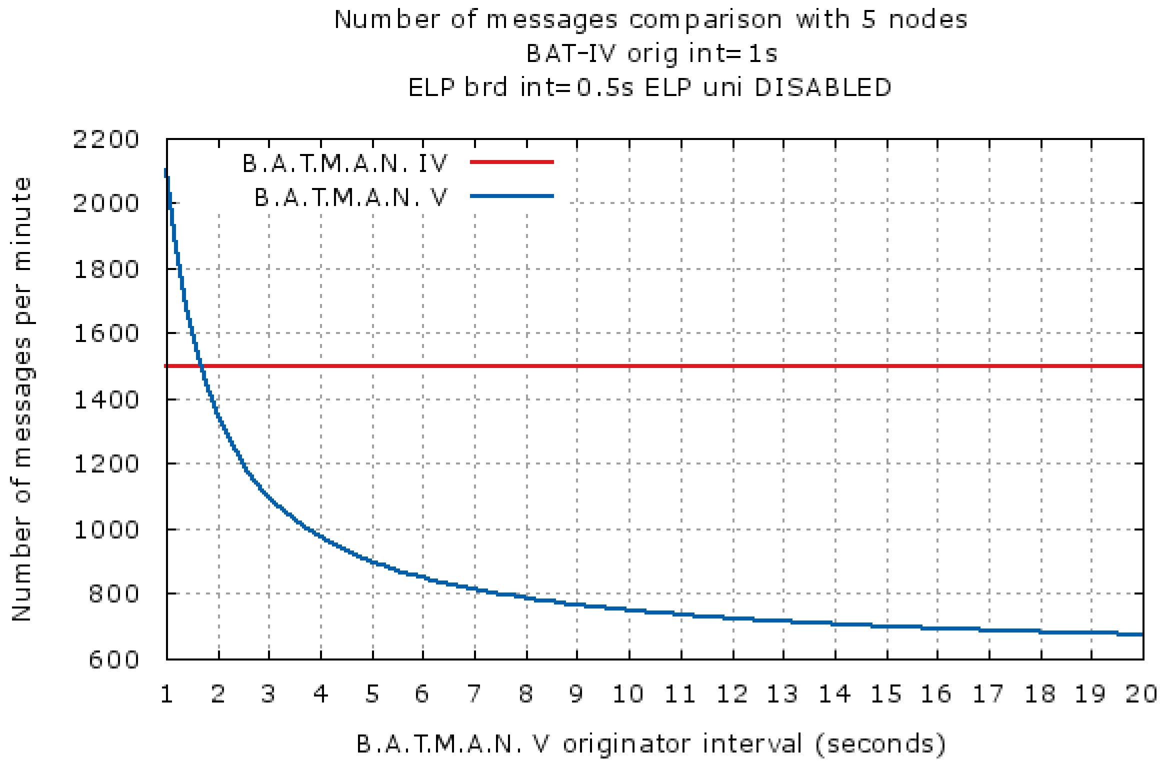

In this chart we can immediately see that, if there were no unicast ELP messages, an originator interval of more than 1.7 (exact value is $1.\overline{6}$) seconds would have been enough to make B.A.T.M.A.N. V perform better then its previous version.

- **0.5s**

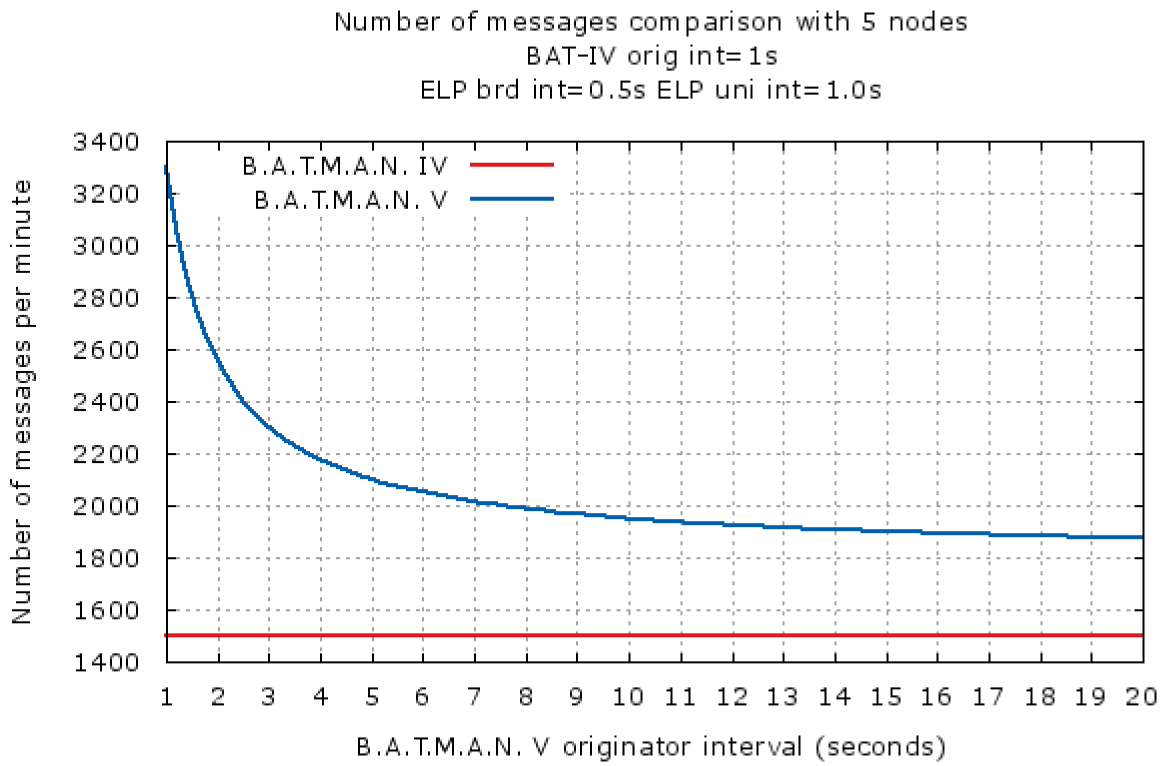

- **1s**

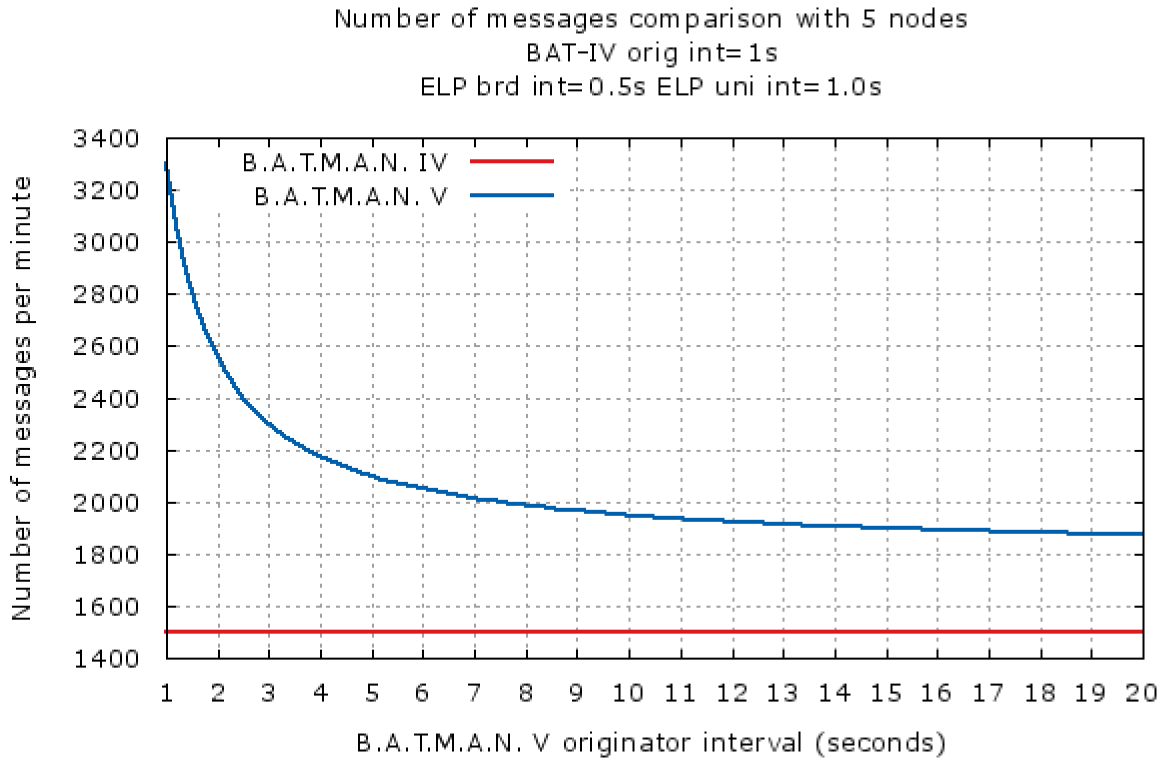

- **2s**

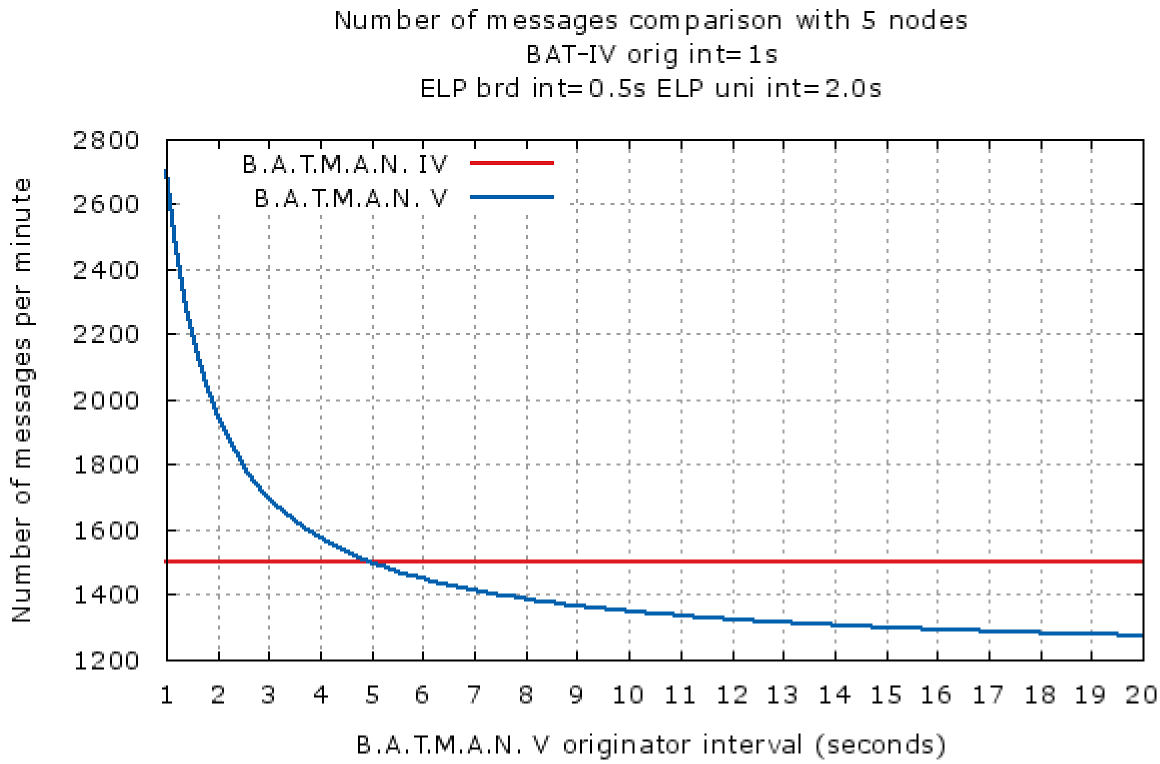

With a unicast ELP interval of 2 seconds we are able, for the first time, to find a reasonable value for the B.A.T.M.A.N. V originator interval which would allow it to generate less overhead than version IV. The result tells us that with an **interval greater than 5 seconds**, B.A.T.M.A.N. V would generate less messages.

- **5s**

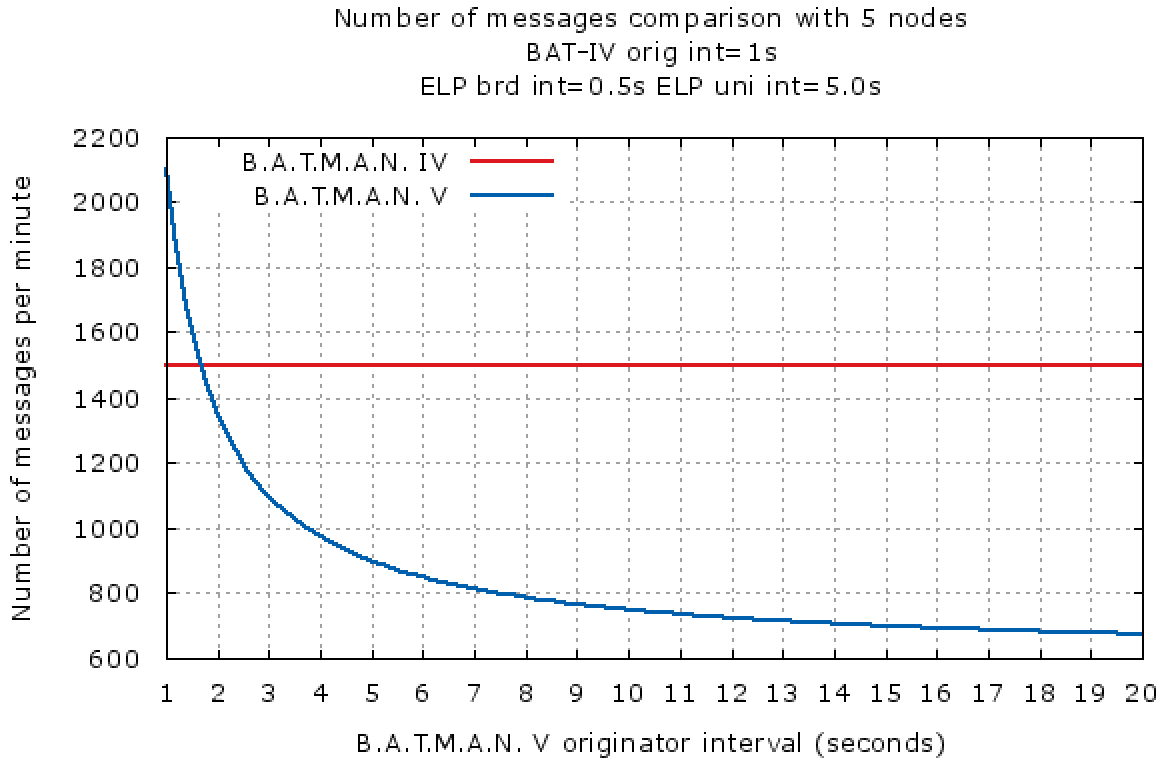

We try to rise the unicast ELP interval more to see how this would affect the “minimum effective originator interval” and as we can see from the chart above, with a unicast ELP interval of 5 seconds, the behaviour of B.A.T.M.A.N. V becomes close to having it disabled.

These results confirm that the unicast ELP packets have an high impact on the whole “protocol pollution”. This suggests us that we have to reduce the overhead given by this component if we want to efficiently increase the performances with respect to B.A.T.M.A.N. IV.

Possible ideas:

1. Increase the unicast ELP interval to 5 seconds (throughput should hopefully not vary very so frequently)
2. Remember the time when the last unicast packet has been transmitted to a given neighbour and send a unicast ELP message to it if and only if this happened more than X seconds ago.
3. Combine the two above and use 5 seconds as value for X.

Considerations:

1. Cost comparison has been done by measuring the “number of packets”, but this may not be realistic since the unicast ELP messages have a cost in terms of airtime which is much lower than a broadcast packets. Moving to the “airtime” metric would allow us to define better thresholds for the unicast ELP interval.
2. Computations of packet rates on a multi-hop topology are needed, because in B.A.T.M.A.N. V ELP messages are not forwarded while OGMs are. This would add another “component” when counting the number of packets traversing a selected area given by the forwarded OGMs.