

WirelessBattleMesh v10

June 5 - June 11, 2017
Museum of Folk Life & Art, Vienna, Austria, Planet Earth
<http://battlemesh.org/BattleMeshV10>



What started 2009 as a “tournament with social character” has over the years become the meeting point of the global community of wireless mesh network developers and activists.

More than 10 local teams from communities all over Europe have organised the event so far - each year a different one. This tradition allows us to share the organisational load and creates an influx of new ideas for every iteration.

In 2017, WBM v10 was hosted by people from a community network project in Vienna, called Funkfeuer, supported by Metalab (open hackerspace) and MinimalKitchen (local food-coop). What was unusual this time was the venue: The Austrian Museum of Folk Life And Folk Art.

Sponsor contributions allowed us to keep the event free of charge, provide free lunch and offer travel scholarships.

This document provides an overview of the event, with budget figures and meeting minutes. Some pictures and a detailed technical report by the testbed team are also included.

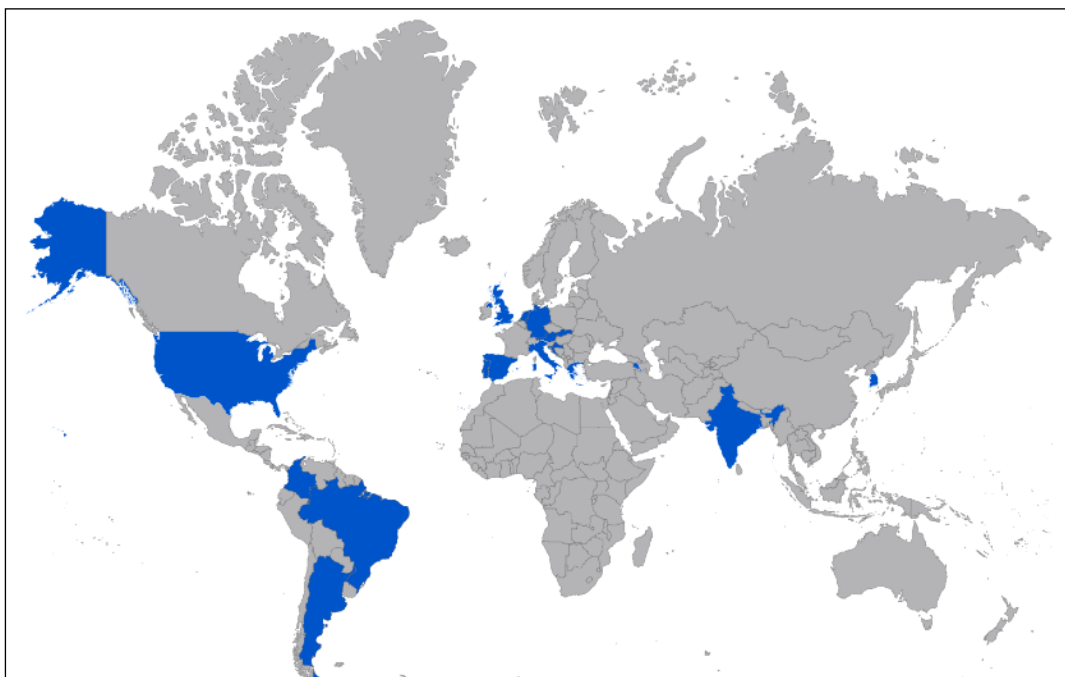
Table Of Contents

1. Participants	2
2. Event Schedule	3
3. List of Talks	4
4. Funding Report	6
5. Feedback from Participants	7
Appendix A: Meeting Minutes	8
Appendix B: Pictures	17
Appendix C: Testbed Results	19

1. Participants

This time, the participants list included 73 People from 16 Nations - which is quite nice - but we want to increase it further in the events to come. Financial support from NGOs for travel scholarships is crucial for that - we want to stay away from commercial sponsoring as much as we can.

This time we even had our friends from Columbia coming over and presenting their work in Spanish - while it was *live translated* to English (thx Nico). During the “future of battlemesh” plenary discussion the idea to try to organise Battlemesh events on other continents in addition to the European one was brought forward. We hope that this idea is picked up soon by our colleagues overseas.



2. Event Schedule

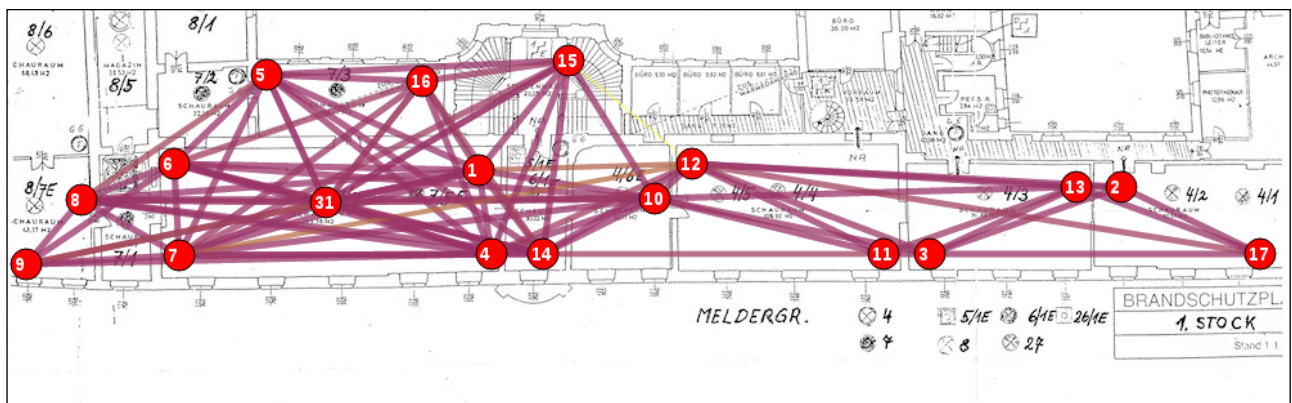
WBMv10 started with a welcome party at <https://metalab.at>, a well known open hacker space in the EU area. During the following 5 days, a total of 27 talks have been held, recorded and live streamed from the main hall of the museum. Thanks to Ryan Taylor a complete public archive is available on youtube:

https://www.youtube.com/watch?v=vPKVeIQk1_0&list=PL3bvPCw5QCLJ-VJPamVeQx-UPNBVyaopj

W23	Mon 5	Tue 6	Wed 7	Thu 8	Fri 9	Sat 10
all-day						
11:00		Morning SYNC - Main Room	Morning SYNC - Main Room	Morning SYNC - Main Room	Morning SYNC - Main Room	Morning SYNC - Main Room
12:00					12:00 Wireless Battle Mesh Future - Main Room	12:00 Workshops - Main Room
13:00		13:00 Community Supported Lunch (Tuesday) - Main Room	13:00 Community Supported Lunch (Wednesday) - Main Room	13:00 Community Supported Lunch (Thursday) - Main Room	13:00 Community Supported Lunch (Friday) - Main Room	13:00 Community Driven Lunch (Saturday) - Main Room
14:00						14:00 IEa, We Scant - clauz Main Room
15:00		14:30 What's up with LEDE? - Felix Fietkau Main Room	14:30 Signed Autoupdates for your Zoo of Embedded Devices done right - B...		15:00 Sovereign Digital Identity - Marius Sabadello	Running Python-based Experiments...
16:00	16:00 Weekly Funkfeuer Community Meeting at Metalab - Main Room	15:30 Olueo - Linus Lüßing Main Room	15:30 Freifunk-Open-MPPT - Elektra Main Room	15:30 Consumer Grade Home Wi-Fi Meshes - apenwarr	GOTHAM - Minsu Kim Main Room	15:30 Social impact of public hotspots - Andi Bräu
17:00		Fiber Deployments Using OpenWRT - Main Room	16:30 BLUECOM+ - Connecting Humans and Systems at Remote Ocean Areas usi...	16:30 Scaling Layer 2 Mesh Protocols - Linus Lüßing	DisasterRadio - Juul Main Room	Leviathan Peer to Peer Network - M...
18:00	18:00 Warmup Party	18:00 ALhea - Jehan Tremback Main Room	17:30 The netCommons project, what are we doing? - Leonardo	17:30 Firmware Lockdown - Updates 2017 - Simon Wunderlich	Funkfeuer Vienna's v842 Project ...	Using Mesh for Rural Education in I...
19:00			18:30 WBMv10: grassroots event turning 10 - Marek Lindner	18:30 Librerouter.org - Main Room	16:30 Battlemesh v10 Meets ISPA v20 - Main Room	17:30 Battlemesh Community Meetup - Main Room
20:00		20:00 Joint Dinner on Tuesday - Main Room	20:00 Joint Dinner on Wednesday - Main Room	19:30 Minstrel-Blues - A Joint Rate and Power Controller for IEEE802.11 devices - Thomas Hähn		20:00 Joint Dinner (Saturday) - Main Room
21:00				21:00 Joint Dinner on Thursday - Main Room		
22:00						

In parallel to the talk schedule, a committed team set up an experimental wireless mesh network on site to test current routing protocol implementations and run performance measurements. The detailed report is included in the appendix of this document and can also be found in the public battlemesh repository:

<https://github.com/battlemesh/battlemeshv10-testbed/raw/master/report/main.pdf>



3. List of Talks

Project And Community News

“What's up with LEDE? A review of the first year of our OpenWrt community reboot”
by Felix Fietkau

“WBMv10: grassroots event turning 10 Where do we go from here?”
by Marek Lindner

“Firmware Lockdown Updates 2017”
by Simon Wunderlich

“The netCommons project, what are we doing?”
by Leonardo

“Social Impact of Public Hotspots”
by Andi Bräu

“Using Mesh for Rural Education in India”
by Senthilkumar M

“Funkfeuer Vienna's v642 Project A Journey Towards IPv6 and OLSRv2”
by David Hopfmüller

Technical Talks

“Gluon A Modular Framework for Your Wireless Mesh Community”
by Linus Lüßing and Matthias Schiffer

“Signed Autoupdates for your Zoo of Embedded Devices done right How to avoid work and lost nodes”
by Bastian Bittorf

“Freifunk-Open-MPPT Open hardware and software solar regulator for autonomous nodes”
by Elektra

“Scaling Layer 2 Mesh Protocols 1000 nodes and beyond with Gluon and batman-adv”
by Linus Lüßing

“IEs, We Scan! Using proprietary beacon extensions to facilitate wireless link building”
by Clauz

“Consumer Grade Home Wi-Fi Meshes“
by Avery

“A case introduction of building wireless mesh network testbed in university”
by Minsu Kim

“BLUECOM+ Connecting Humans and Systems at Remote Ocean Areas using Cost-effective Broadband Communications”

by Filipe B. Teixeira

“Althea Incentivized Mesh”

by Jehan Tremback

“KadNode: A p2p DNS system based on a DHT”

by mwarning

“Disaster Radio An off-grid low-bandwidth long-range emergency mesh”

by juul

“Replicating Battlemesh on open wireless testbeds”

by Charalampos Manolidis

“Minstrel-Blues - A Joint Rate and Power Controller for IEEE802.11 devices Status of implementation, usage and performance”

by Thomas Hühn

“Social and Motivational Aspects - What Makes Community Networks Thrive?”

by Lusy and Leonardo

“Running Python-based Experiments on Routers with Seattle”

by Albert

“Digital Sovereign Identity”

by Markus Sabadello

Software Releases

During the event, two software releases took place:

Gluon v2017.1 10.06.2017

<https://github.com/freifunk-gluon/gluon/releases/tag/v2017.1>

LEDE v17.01.2 12.06.2017

<https://lede-project.org/releases/17.01/notes-17.01.2>

LEDE and Gluon are currently the two most important FLOSS projects for mesh-networking :D

4. Funding Report

With help from ISOC and INESC we could financially support people from Columbia and Brazil to offset their long distance travel costs. Also, thanks to a number of smaller local sponsors (Freewave, ISPA, IPA) we could provide free lunch on 5 days and offer caffeinated beverages at a reduced price throughout the whole week.

The remaining costs have been graciously covered by the Funkfeuer Wien Association - resulting in balanced budget in the end. All financial handling was done by Paul Fuxjaeger - a trained electrical engineer with a strong aversion of accounting work - but looking back he thinks that the cause was definitely worth it :)

Sponsors	Amount (EUR)
Internet Society - https://www.internetsociety.org	5,829.24
INSEC Tec - https://www.inesctec.pt/	1,423.54
FunkFeuer Wien Association - https://www.funkfeuer.at	1,085.00
Freewave GmbH - https://www.freewave.at	1,000.00
Internet Service Providers Austria - https://www.ispa.at/	450.00
Internet Privatstiftung Austria - https://www.nic.at/de/das-unternehmen/ipa	300.00
Cash donations from participants	51.92
Total Event Budget	10139.70

Costs	Amount (EUR)
Location Rent	3850.00
Museum Security Guard 24/7	670.00
Travel Scholarships	2823.95
Accomodation Scholarships	1165.81
Foodcop Expenses and Beverages	1561.47
Other Expenses	66.47
Total Event Costs	10137.70

5. Feedback from Participants

Anonymously collected notes on the WBMv10 feedback page:

Like:

####

- lots of creative space was nice, courtyard and nearby park was cool (paul) +1 +1+1++1
- community supported lunch worked well (paul) +1+1
- talks were all high quality, spam level was close to zero (paul)
- most talks were properly streamed and recorded. Thanks to ryan we also have a fully documented backup of v10 on youtube (paul) <https://www.youtube.com/channel/UCxfh-2aOR5hZUjxJLQ2CIHw> +1
- 20+ *identical* devices (wdr4300 in our case) readily available for testbed team (paul)+1
- awesome location
- very good deal for accomodation
- extremely fast deployment of testbed (compared to previous editions); and its coordination was spontaneous and natural, agree?
- the possibility to hang out at the metalab after hours (and having a local community support the event!)
- meeting everybody of course :)
- talk from bogotamesh
- the good stuff in tetra pack
- the orga team was great!!1010101!!!
- crazy dancing (or shaking bones) moment

Dislike:

#####

- talk schedule was prepared too late, changed a lot and was published poorly, sorry for that! (paul) +1
- some dinner plans were done poorly (paul)
- internet access had some outages (paul) +1 (although it was not that terribly bad to not hang in front of the screen all the time^^)
- signal/noise situation in the main room was suboptimal (paul) +1 (although it has been worse before, it was not so bad) / Having talks in it's own room (Maribor) is disliked by some people. Hard to find a solution everyone is happy with. / I liked the idea of almost everything happening in the same place
- hostel reservation was done too late (paul)
- endorsements were done poorly (paul) +1
- internet connection in the venue was flacky when all the participants were there. Well, this is just an optimization, but maybe in the future we could try to make this slightly better - we expect flawless WiFi connection just like our flawless mesh protocols :D
- pauls self-criticism

Other comments:

#####

- who noticed that a person working in facebook's connectivity lab was also participating? :-D oha I did. Maybe they are now inspired? :) why does this actually matter?
- who noticed that a lot of strange people attended the battlemesh?

Appendix A: Meeting Minutes

Future of Battlemesh - Plenary Discussion

It is day 2 and the testbed is running this year (please give a round of applause from home as well)

Setting the issue:

- * there is a lot of frustration over the testbeds - some folks who have stayed away because they were offended by some way in which the tests were conducted
- * lack of preparation
- * People preparing the testbeds get frustrated because of criticism and complaints over the testbeds targeted at them.
- * Are we trying to re-introduce people or avoid further frustrations?
- * A unique part of the event is the communication with the people. It helps meeting people that are involved in every aspect of the Mesh (firmware, communities, hardware etc)
- * Is there a reason to call it battle?
- * Is the full week duration the best format?
- * for the tests we need the week
- * If the event is shorter it might attract people that have less time (though they can participate anyway in the 1 week format)
- * People arriving in the middle of the week might feel lost
- * Start on Friday get the work done more structured and then leave a few days for socialisation
- * Does this have to be a separate event?
- * Core team is stepping down

What should we do?

- * Nico: Get in contact with communities that are most in need of your expertise in order to connect more with the social aspect and understand users.
- * Set a challenge for the event every year (in a community of example). Everyone can know it in advanced so they can think and prepare ahead of time.
- * Let's make skirts with the event logo
- * Participants coming to the event with issues each person/community has.
- * when the pitching of problems has been done then the community can pick what interests them to work on.
- * Try to set a structure to the testbeds in order to get them working for example having several testbeds for different configurations where it makes sense to compare protocol A to protocol B instead of having only one.
- * use it as a tangible demonstration of features rather than comparison.
- * Work out the testbed frustrations during the event.
- * pick what to check every year depending on what is interesting.
- * Compare families of firmware instead of individual ones, with the goal of improving.
- * Try to be more inclusive (for example gender balanced)
- * the friendly competition of the WBM is not about winning the battle but about the learning experience along the way: There is no reason to have a winner.
Maybe let that be visible in the declarations and website?
- * Simplify the test: From last years we learned that what we are trying to do is very complicated.
- * Let the testing be organic and optional. If the testing is not planned there will be less pressure and anything tested will be a plus.
- * tutorials for new people
- * social activities/gamification:
 - * (break the testbed!)introduce bugs and let 2 teams solve the issues
 - * flood the network with (virtual) 1000 nodes joining

- * soldering against Bastian
- * Configure a Mesh network vs time.
- * pre-events might be able to separate social from technical aspect of the event

Proposal for games (social event) for two teams
/* terminal on the beamer; cheering; laughing */

* game 1: mesh-bug-hunting-combo

- the jury places like 10 or 20 bugs in a working wbm-testbad and the teams have 15min or 30min to pin them down

* game 2: init-the-mesh-combo

which team is faster on bootstrapping a working mesh (batman, olsrv1, olsrv2, etc etc)? goal is to transfer a file from one mesh node to another over the dynamic routed network. google is allowed. teams start with a naked .config and build images only with vanilla packages included.

* game 3 (special): speed-soldering against bastian

how is soldering faster a serial connection on different router modules then bastian?

Who organizes that?

* Who are the core team?: Pinging people in order to keep them motivated. This requires discipline and dedication. Make sure someone is incentivised to do the Mesh. Make sure there is someone to do the next.

* What was their motivation?: Learning how to organise events. See the event happening. No one else was doing it. In the beginning there was a group talking wireless in the CCC and they felt they wanted a technical event for wireless. For the fun of it.

* It is not a hard part but it is a commitment. It is mentoring, counseling and reminding. Perhaps having not just the prior year, but the previous team before that, make themselves available to mentor the next

Can the local team find the next team?

* It might be too much pressure on them.

Maybe if we add a documentation to the mix?

* We have a framework of what has to be done by when and by whom.

Having a core team might help with preserving information and experience to the next events

* Might be a good idea to have someone(s) who can commit themselves to the role for 5 and not necessarily for 10 years

It has always been a responsibility of the local team to help mentoring the next.

Distinguish between of 2 categories

People who have no time

hard to manage

People who are afraid of doing it badly

can be helped to take the responsibility and the exposure

What is the point of this:?

Paul: Desire to keep this tradition alive, and not just let it fade away. What I've experienced with this event is some confusion over the direction - is the focus the testbeds, the social.

If we find a common denominator then the problems can be solved because we know the common goal.

This feels like a tiny version of the CC Congress... I think all we have to do is find another name, actually.

Find our mission!!!

It is not a battle.

We had the discussion of "what is this?" over the years. Every local team makes it their own in a way or another.

There are a couple of days to discuss about it.

Marek: We won't disappear, but we just want to take a step back.

Try Lusy's idea of having the current team ensuring that the next event happens. And let's start it now: Get a destination for the next event.

Declarations:

- * Filipe: Would be free to organize it - not next year, but in 2 years or so.
- * Marek and Simon: We will be here for you.
- * Paul: It was very helpful to have clear separation of competences.

Locations for next event:

- * Jenny: USA - (Oakland California - <https://omnicommons.org> - home of <https://sudoroom.org> ... and also Food Not Bombs...) - there had been conversations in 2015 about a 'Battlemesh West'
- * Gui: Formentera (zoobab: +1)
- * Daniel: Paris, 10 years after Babel being presented that would close the circle <https://www.irif.fr/~jch/software/babel/babel-funkfeuer.pdf>
- * We could even suggest location where there is no one (for radio noise it is better in the countryside), e.g. co-located with <https://camp.hsbg.org> at https://en.wikipedia.org/wiki/Fort_Monostor
- * Weimar?

Questions for the Round Table meeting:

- * Who is stepping up as mentors?
- ** let's not make a 10 years plan.
- ** let's have someone for next year and we see.
- ** we need a mentor not a recon team necessarily.
- ** Paul volunteered for this year.
- ** Maybe the current team mentor the next. And solidify this in the takeover agreement.
- ** Who is helping Paul? We need a volunteer:
- *** Claudio
- *** Filipe
- ** In every battle mesh we can have a meeting at the end to find also mentors.
- * Do we need a recon team?
- ** If the local team is already experienced, we should skip the recon step.
- ** Can be helpful to have a fresh pair of eyes on an area
- ** Having a recon team allows newbies to organize
- ** Helps to focus the local team.
- ** Helps to transfer a lot of knowledge in short period of time.
- ** Locals can see commitment.
- ** Useful to set up a deadline for the locals
- ** Can be expensive to the recon team.
- ** Maybe sponsorships can help the recon team
- ** Money is not really a problem we need to solve - more so it's motivation to organize
- ** Current year organizers 'voluntarily' mentor the next and take the role of 'last man standing' with the help of previous years organizers
- * Formalising Battle Mesh?
- ** Making an association might help with sponsorships
- ** Moving money might be easier, and money could transfer from one year to the next
- ** So far the money have been handled by existing (local) associations.
- ** Sponsors usually are locals that give a bit of money. And for the past few years there was ISOC as well.(they are very enthusiastic about what we're doing)
- ** maybe other organisations can help us make an umbrella.
- ** Getting attached to ISOC would be getting tools from them. They usually don't impose stuff.
- ** Ispa might be helpful.
- * Are we changing the name to something without "battle" in the name?
- ** International summit for Mesh Networks
- ** The discussion can continue in the mailing list and have a vote.
- ** Pau and Gui(?) will initiate the conversation on the list
- ** also change the first line in the website?
- ** Battle FOR the mesh and not OF the mesh so it looks more unified
- ** Filipe: had some problems explaining to potential sponsors why we are 'battle' mesh
- * Are we changing the format for the next events?
- ** Can steer it a bit by changing topics/scope eg, focusing more on the community aspects rather than purely technical

- ** Will influence whether we have a more professional format or say hacker camping
- ** Could change focus to having more of an unconference format
- *** Seems pretty much like an unconference already
- ** Originally there were no talks at all
- ** Event only needs a location and a date - doesn't need a fancy tshirt, doesn't need a talk schedule
- ** outsourcing items in the wiki is a good thing and we are doing it
- ** List on the wiki of tasks that need to be done by local teams and tasks that can be done remotely
- ** Maybe offer a pack of pre-backed items to help the teams.
- ** Make a list of "esencial parts" in order to make the local organiser's life easier.
- ** Have a meeting at the end of every Battlemesh formalized around creating the arrangement for the next year's battlemesh
- ** LEt's have workshops in order to help new people to not get lost
- *** We can have a workshop schedule just in the same way we have a talk schdule
- ** Do we want an un-conference?
- *** it will attract more people to have an agenda a couple of weeks before the event.
- *** For people abroad it is even helpfull to have is monthes in advance
- ** Let's let people know that everyone is aproachable so the new people can find thier way
- ** We can mix them both and have talks as well as unconference time
- ** Paul: Overwhelming feedback has been that the mix of scheduled and ad-hoc activities has been balanced well at this event
- ** There is value in attracting new people otherwise we will age out of this
- * How do we continue supporting an event in the current scale?
- ** Suggestion to make the format leaner - eg, leave accommodations up to the participants
- ** Seems like that is unlikely given that the event continues to grow in # of participants and international participation
- ** Do we want to change the scale of the event?
- ** we can fix the number of people via altering the marketing strategy
- ** The participants have been stabilised for some years now to ~100
- ** Mentoring team that stretches over the 3 previous years, eg for WBM11, mentorship team is from 8, 9 & 10
- * Focus for the next events (experiments, talks, social events?)
- ** May be good to have more educational workshops, eg firmware building 101
- ** Could also go to a remote location with a challenge to eg set up a long distance link
- * Duration of the event?
- ** decided by the organising team
- * Do we want to have a stable location or shift locations?
- ** Worries that would result in inertia
- * Do we want sponsorships? If yes What are the opportunities/possibilities available?
- ** Somehow covered by previous discution.
- ** ISOC would likely be willing to more formally sponsor
- * Do we want to reach out to Academia/Universities?
- ** we could do more outreach to local universities
- *** this alread happens, question is whether to bring in outside academia to the event to make it more educational / formalized
- ** Bringing outside accademia? With no papaers and accademic stuff it will be hard. We need results.
- ** Elektra: prefer a more informal character, not changing to accommodate the stricter needs of academia
- ** Pau: we could just do nothing, and if universities want to get involved they can
- ** Lets not cater for their demands or needs
- ** Let's make the event more easy for newcomers - workshops during the first day(s)

“What Makes Community Networks Thrive?” - Plenary Discussion

* Time and Location: Thursday, 4:20pm, outside in the grass

== Participants ==

- * Lusy (Freifunk Berlin)
- * Leonardo (Ninux, NetCommons)
- * Jenny (People's Open Network/sudo mesh)
- * Reinhard (Tripoli Labs)
- * Claudio (Ninux)
- * Pedro (Guifi)
- * Juul (People's Open Network/sudo mesh)
- * Krishna
- * Peggy (Freifunk Berlin)
- * Phillip (Freifunk Berlin)
- * Rory (Researcher, nyc)
- * Greg (aspiring, Atlanta, Georgia, USA)
- * Nefeli (Freifunk)_

== Notes ==

- * Telecom sans frontiers - French org working in Africa
- * UNHCR

* Reinhard, Greece - having lots of problems getting people involved - have tried for 8 years, tried outreach to a local technical school, but no one chose to participate in any way. Tried to combine with a summer camp - that's upcoming. Visited AWMN in Athens and will be inviting them to come and help. Also inviting folks to come and help at Battlemesh. No special tasks to fulfil.

* Claudio, Rome (Ninux) - had some years of big growth but has become static. Good communication and technical skills, but

don't have a good automatic system - very manual - growth changes a lot.

* Pedro, Guifi.net - very large and heterogeneous, lots of users. See companies are helping to make it more sustainable... but growing stagnant in volunteer energy. Professional part is growing, but unfortunately the grassroots element is diminishing.

* Krishna (Germany) - not part of a community network yet. Trying to understand what motivates us. Like the technology of mesh,. Like the idea of a community coming together, nonmonetary motivation, collaboration and content creation. Want to know why you do

* Jenny - / <https://peoplesopen.net/sudomesh>

* Peggy (Freifunk Berlin) - Refugee camps in Berlin, just came back from traveling throughout Brazil talking with different communities and also with an official org in Brazil about the issues around access and education there. Need internet to support agricultural business for example - connecting isolated areas and farmers to the cities where they sell their products. Huge class disparity in who gets access - Brazil is one example of a phenomenon that's happening all over the world.

** Reinhard - we do have the same situation in Greece, where cost is 4x that of Germany, income is 1/4 of Germans.

*** Poor people if at all have a prepaid sim card. Absolutely forbidden to share internet - by contract and by law. However there are workarounds, a book has been written - [blog url]

* Lusy (Freifunk Berlin) - wrote a paper last year on participating in freifunk. would be interesting to ask other communities. there were already differences withing freifunk since it is locally organized. even more so when we e.g. compare ninux with guifi with sudomesh etc. every project functions in a different. people have somewhat different motivations and ways to engage newcomers and how to organize themselves. kinda a follow-up on a larger scale.

* Leonardo: from ninux in florence. working at university in Trento where i'm working ont his project i was talking about (<http://www.netcommons.eu>). we had a ninux meeting in florence and there were 20-30 people from all over italy and asked them what their motivation for participation in ninux was. half was technical, half was digital divide, 1/3 social. in italy we have 80% coverage of broadband. lowest in europe apart from greece. europe is investing a lot in fiber. you may think that this will solve the problem but with fiber you get more bandwidth but the service is more costly (in italy). if you can pay for it you can have fiber so it brings more bandwidth to the people

who can already afford bandwidth. Trying to build a grassroots alternative to the system that's already in place. Believed that we could build alternative services that could replace Google & Facebook and failed. Cannot fill it with services that people won't use. Globally there is a lot of attention and initiatives around community networks around the world. IGF (?)... lot of bottom-up experiments and attention to them... but that is frequently mismatched.

* Marc (juul) - Peoples Open Network in Oakland - learning by building our own stuff these past few years. Part of motivation is that our ISPs suck so badly - options are Comcast or AT&T, many areas only have one of those options. Intiatlly when stuff went down we thought it was us - but usually it was actually the ISP that was down. Even in the Bay Area, fiber is unavailable or unaffordable (\$2200/month aside from the cost of buildout). Unaffordable for our community space to pay for it. Bandiwidth is pretty cheap (\$150 for a gigabit). Had a lot of trouble just getting bandwidth into our network. Just getting something in place that would be resilient is a motivation... working on a low bandwidth solar-powered emergency comms network

* Philipp - Freifunk: Since 2015 I'm in community networks. Motivation changing over time, currently very diverse motivations... one driver is connectivity, refugee crisis of the last two years, lots of people in need of communications infrastructure. Lots of people got involved recently because of this. Most of the city areas have proper coverage with DSL or ? - in the past the project was more of a nerd playground, philosophical inclination toward decentralization. If the project wants to grow, a purely play/philosophical-driven network can only grow so much. Interesting to bring together different ideas and build a larger cooperation around

* Greg (Atlanta GA) looking to start a community. Learning from coming here, what works, best practices, etc

* Rory (New York): Observer, fascinated by mesh networks. social, political issues, when where people do this. I'm writing a research project that may turn into a technical history of the socio-political aspects of mesh networks.

== Conversation ==

* Pedro - Curious to hear more about Ninux' attempt to run services on the network... started researching what could be done. First thought - put content on the network - but adds an additional layer of complexity (the content must be maintained). For communications - after one year, people love using the network to chat. We use RocketChat . Also JitsiMeet for videoconferencing <https://meet.jit.si/>.

* Claudio - Had a Jabber serving running - was very successful until it stopped being maintained.

* Philip - Should combine the network with services in the network. Can really benefit from making a very clear separation between the network and the services - as in Guifi.net - which makes some things very easy. Internet is not a default part of the network. In Frefunk, internet is the default - creates problems with the technical implementation, but is also a good motivation, because if you participate in the network you can get internet. If you make a clear separation, it's better.

* Reinhard - We called our network i4free - but people said it's wrong, it's not free - but for the people who need it, it is free.

** AWMN died - no more money flow. That was done initially with lots of money - but the money and people went away. After a peak growth phase, need to reexamine the motivation behind building and sustaining the network

* Pedro: On Guifi, we run several services. Chat is important for mentoring new people

** Krishna: But 10 users out of 30,000 nodes? People don't want to use alternative services, so it makes more sense to just provide internet, with services available as a backup in the event the internet is unavailable

* Leo: Those networks that are large grow because they provide internet - the other primary motivation is sovereignty, so providing internet is antithetical to that

* Krishna: I came to battlemesh because I was interested in the idea of "To bring Internet to people that don't have Internet"

* Peggy: It isn't up to us to determine what people should use the network for. They will do with it what they need - eg, searching for markets in which to sell their wares. The children go to town frequently just to access the internet. What they're doing is up to them. I met an indigenous woman who uses internet primarily to run a karaoke app. It's important that we do bring internet to people, because if we don't do it, others will do it (like Internet.org from Facebook) very badly. People need to organize in their own communities, reach out if they need help.

- * Leonardo: If we just provide internet access, eventually some telecom operator sees a market need and comes and takes control of it, without resistance - because they can do a better job.
- * Nefeli: People use services on the network to do "local stuff" - share information about local events and sales at local markets
- (...)
- * Marc: Use services as a way to store local information that is secret - eg, underground parties - encouraging people to switch to alternative services (eg Patchwork in place of Facebook)
- (...)
- * Claudio: we lost the facebook train, we could still be able to catch the chat train
- * Philip: Formal agreement of how to do this corporation
- * Pedro: Formal agreement is very well defined, probably not very well documented. But this does not help about the "healthy community topic". I'm happy this is happening to battlemesh too, we are users, and the organization team is low. How to have more active people?
- (...)
- * Krishna: how to

== Links ==

- * <http://freedigitalterritories.pexlab.space/> - Peggy's blog, with interviews from her trip through Brazil researching the sociopoliticotechnical landscape of current and potential community networks
- * <http://artigo19.org/wp-content/blogs.dir/24/files/2017/01/Como-Montar-e-Regularizar-um-Provedor-Comunit%C3%A1rio1.pdf> - Booklet for setting up community internet providers in Brazil. Technical and legal practical guide. (in Portugese)
- * https://commons.thefnf.org/index.php/Network_Commons_License - Efforts toward creating a Network Commons License. Collaboration between Free Network Foundation, sudo mesh, Guifi, Freifunk, Funkfeuer(?), Altermundi, & Ninux

Ad-Hoc Workshops

Workshop Mode:

- one Moderator to "start" the workshop
- document in etherpad during workshop, copy link to wiki
- remote participation: etherpad (maybe stream too?)

Higher Layer Topics:

“Social and Motivational Aspects - What Makes Community Networks Thrive?”

- * Moderated by: Lusy and Leonardo?
 - * Time and Location: Thursday, 4:20pm, outside in the grass
 - * Interested participants: Jenny
- <https://etherpad.funkfeuer.at/p/wbm-social>

“Firmware Show and Tell - let everybody know about new features of your favourite firmware”

- * Moderated by:
- * Time and Location:
- * Interested participants: mwarning

“About APIs - Mgmt, NodeDBs, Visualizing - Short overview and update on active projects”

- Let's create an overview of currently running projects in this area
- * Moderated by:
 - * Time and Location:
 - * Interested participants:

Routing Layer Topics:

"IPv6 in layer 3 mesh networks"

- * Moderated by:

- * Time and Location:
- * Interested participants:

"Prototyping routing protocols in node.js"

Source code: <https://github.com/booo/node-distancevector>

Requirements: Laptop, basic programming (copy/paste) skills

Let's learn how to implement a simple routing protocol from scratch in a few lines of javascript.

Please flip a bit if there is any interest: 0000 0000 0000 000

* Moderated by: Philipp (Freifunk Berlin)

* Time and Location:

* Interested participants:

Lower Layers Topics:

"Sharing our experiences of using DFS channels in a mesh configuration"

* Moderated by:

* Time and Location: Sat 12:00pm

* Interested participants: Paul, Simon, txt.file, mwarning

"Switching from IBSS to 802.11s (w/o forwarding): Pros & Cons"

* Moderated by:

* Time and Location: Sat 12:30pm

* Interested participants: Simon, mwarning

"Current state of 802.11s mode in open source drivers"

* Moderated by:

* Time and Location:

* Interested participants: Paul, txt.file

"Future Outlook on Available Radios+Drivers with Meshing Capabilities - Life After ath9k"

* Moderated by:

* Time and Location:

* Interested participants: ?? no one knows anything or no one is interested?

Experiment Ideas

MAC Layer:

<http://ml.ninux.org/pipermail/battlemesh/2017-March/005410.html>

- is it optimal to use 802.11s(norouting) instead of IBSS?
- what's the potential of per packet power control?

Airtime fairness for ad-hoc

Does airtime fairness scheduling make sense in mesh?

<http://ml.ninux.org/pipermail/battlemesh/2016-November/005239.html>

Airtime fairness seems to be available for ad-hoc (their focus was on the AP mode, which is bringing a massive improvement for latency according to the slides):

http://www.linuxplumbersconf.com/2016/ocw//system/presentations/3963/original/linuxplumbers_wifi_latency-3Nov.pdf

That would be great to make a test before/after the patch just for the ad-hoc case.

Network Layer:

<http://ml.ninux.org/pipermail/battlemesh/2017-March/005412.html>

- what's currently the optimal throughput metric to use?
- is radio-diversity routing (bmx7, Babel-Z) worth it, and in what topologies? (you mean if testing routing protocols that still hop over the same channel still make sense -> no, because of poor performance due to the half-duplex nature of the shared spectrum)

- is RTT-based routing (Babel-RTT) useful in meshes? (It's known to be useful in overlay networks, which is what it has been designed for.)
- is source-specific routing useful for the mesh community?
- can we use ToS for something?
- is there any sense in developing autoconfiguration protocols, or is DHCPv6-PD good enough?

Test firmwares instead of routing protocols

<http://ml.ninux.org/pipermail/battsed> on the OpenWRT Blemesh/2016-December/005263.html

Example: on day 1 of battlemesh, qMp devs flash everything with qMp firm, everything works out of the box since they care about that every day of the year :) qMp devs can explain a bit the features, routing architecture, etc and it runs bmx7 so Axel can even run tests on his protocol that day

Next day, gluon devs flash a gluon community firmware (could be something done specifically for the event, or just extending an actual community). again, everything works out of the box, and has a sensible configuration of batman-adv

gluon devs can explain features, config wizard, etc
batman-adv devs can run tests on their protocol that day

next day, libremesh devs flash libremesh... rinse and repeat

- Gluon
- libremesh
- qmp
- wibed
- vanilla LEDE/OpenWRT

BATMAN filter

```
tcpdump -ni fastd-mesh0 ether proto 0x4305 and ether[14]!=0x40 and ether[14]!=0x41 and ether[14]!=0x42 and ether[14]!=0x01
```

```
0x4305 = batman-adv ethertype
0x40 = unicast
0x41 = unicast-frag
0x42 = unicast-4addr
0x01 = bcast
```

Stress Tests

<http://ml.ninux.org/pipermail/battlemesh/2017-March/005380.html>

I have been stress testing multicast (after getting ATF to work on unicast this past year), and ended up rolling a few tools that let me abuse it in the case of meshy routing protocols.

Last night's effort is called "rtod - routing tables of death".

The initial code is here:


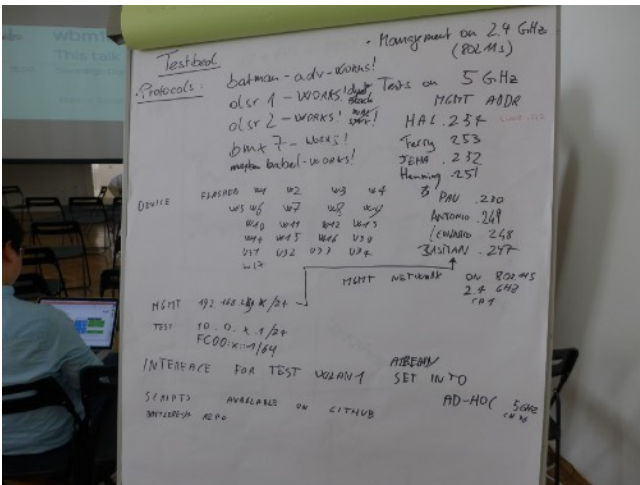
<https://github.com/dtaht/rtod>

Would be interesting to see if bmx7 (with security extensions [1]) is able to minimize the impact of such kind of attack. Of course the network flooding will still be a problem but if the routing protocol can survive to it the attack impact would reduce drastically.

[1] <http://bmx6.net/documents/30>

"Mesh" is a big thing in eero and google this year... why not test that too. (802.11s)

Appendix B: Pictures





Report on the testbed results of the BattleMesh V10

Bastian BITTORF
Alessandro GNAGNI
Leonardo MACCARI
Giacchino MAZZUCCO
Claudio PISA
Jehan TREMBACK

Abstract

The Wireless Battle of the Mesh is an yearly event that brings together people from across the world to test and compare the performance of different routing protocols for ad-hoc and mesh networks, like Babel, B.A.T.M.A.N., BMX6, OLSR, 802.11s. Every year the community gathers and set-up a testbed on which the protocols are run, developed, debugged and tested, and some performance measures are extracted. While the initial spirit of the event was to set-up a competition between the protocols (as the name suggests), with time it changed into a moment of exchange of experience, collective development of innovations in the field of mesh networks and wireless open source networking software. This document reports on the experimental results of the tenth edition of the Battle of The Mesh event.

June 20, 2017

1 BattleMesh v10

The Battle of the Mesh (WBM), as the official website says:

It is a tournament with a social character. If you are a mesh networking enthusiast, community networking activist, or have an interest in mesh networks you might want to check this out!

The goal of the WirelessBattleMesh events is to set-up hands-on testbed for each available mesh routing protocol with a standard test procedure for the different mesh networks. During the different WBM events, similar hardware and software configuration will be used based on the OpenWRT BoardSupportPackage and packages for each protocol implementation. The WBM events are also a great opportunity to develop testing tools for PHY/MAC radio layers (drivers, scripts and PHY analyzers).

WBM is now at its tenth edition, and is organized by a motivated and large group of people (approximately 80-100 participants in the whole week in the last editions from 2-3 continents).

2 The Testbed

This year, the testbed was realized with 17 TP-Link WDR4300 routers, equipped with two wireless interfaces (operating in the 2.4 and 5.0 GHz bands), and 5 Ubiquiti Unifi AC Pro. The latter were used to perform local testing and support to the tests, but did not participate to the routing, since they are equipped with a different chipset. The nodes were configured to set-up two independent networks, a “management” network, running on the 2.4GHz, and a “testing” network, running on the 5 GHz. The management network was configured with the IEEE 802.11s protocol, and was used only to access the nodes and perform tasks. The testing network was made with interfaces configured in ad-hoc mode and, for each test, was using a specific routing protocol.

Figure 1 reports the topology of the network as exported by one of the network node, when running the OLSRv1 protocol. This topology was used to understand and roughly guess the property of the network. The underlying image is the map of *Volkskundemuseum Wien*¹ that hosted the event, roughly, the distance from node 9 to node 17 is about 75m, and the width of the main room is about 8m (the main room is the room where nodes 1,4,7,6,31 are placed and where the conference took place).

Routers numbered from 1 to 17 are the TP-Link, while router 31 is an Ubiquiti router that is used in this case to extract the network topology. In the real tests this router does not participate to the network functioning. The transparency of the links represents the badness of the link quality as reported by the OLSRv1 protocol, using the ETX metric. A solid link means $ETX = 1$ (good link), while a transparent link means $ETX > 1$ (the highest, the worse).

Note that the topology in fig. 1 does not necessarily represent the topology used by other routing protocols, but gives an approximate idea of what links are directly

¹ The Austrian Museum of Folk Life And Folk Art, Laudongasse 15-19, Vienna, Austria

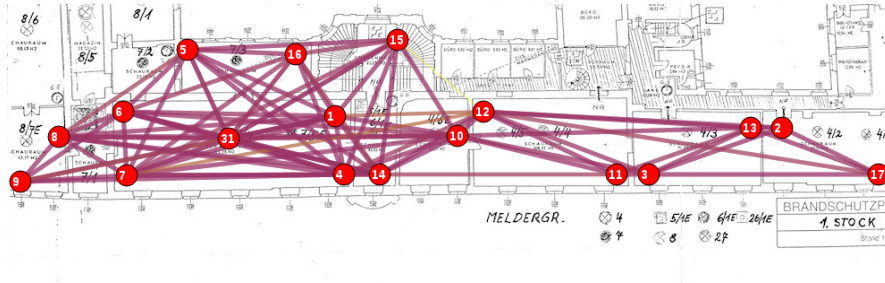


Fig. 1: Network topology, as exported from the OLSRv1 protocol

connected. Based on this topology, exported in the netJSON format², we calculated the weighted shortest path between any couple of node, using the networkX python graph library. We repeated the process twice, with two different transmission power level for the testing network. With the transmission power set to 17dBm the network was considered too dense, and thus of little interest for the routing function. Therefore the power was lowered to 10dBm, and produced the distribution of path costs reported in fig. 2. Each entry is computed as follows, given the network graph and a couple of nodes (A, B), the shortest path between A and B is computed via networkX and then the sum of the ETX values for the path is done³. We selected four nodes to perform the tests, node 8, 9, 17, and 2, that are at the extreme ends of the topology. fig. 2 reports the corresponding shortest path costs in the ranking.

2.1 The Protocols and the Experiments

Several protocols were tested during the WBM, not all the protocols (or their variants) have been tested on all the configuration. Table 1 reports the list of the protocols, a brief description, and the link to the source code.

Four out of six days that make the WBM were devoted to set-up the testbed, and only the last two were dedicated to the testing itself. The procedure of set-up and testing is error prone due to a vast set of reasons that range from the need to use a recent OpenWRT/LEDE version, the potential incompatibility with the last version of the protocols, their configuration, and last, but not least, the eventual bugs that are found while testing and need to be fixed. This is a very important part of the WBM, possibly the most important under the technical point of view (the social side of the WBM as equally important). During the tests the developers of the protocols (that are generally present at the event) test new features, compare different strategies and inevitably stumble upon unknown bugs in their protocols. This process is vital for them to improve their software and stabilize their code, and is arguably even more important than the results of the tests themselves.

² A generic format for exporting a network topology, see <http://netjson.org>

³ For the code used for this task, see https://github.com/battlemesh/battlemeshv10-testbed/blob/master/tests/failure_recovery/parser_scripts/parse_json.py




Fig. 2: The ranked list of the ETX weight of all the shortest paths in the network.

Four set of experiments were planned, and three were fully performed:

- Ping test: a session of 100 pings from node 8 to 17, 8 to 2, 9 to 17, and 9 to 2 were performed to measure the loss and the delay distribution. This test was repeated only once.
- iperf test: a batch of 10 TCP iperf sessions were run (10s each) between node 8 and 17 and from node 9 to node 2. This experiment was repeated with two more variants, one in which node 14 and 10 were running another iperf with a 5Mbps limit, and a second one in which 4 5Mbps session was running between other nodes. These added sessions were not recorded but where used to increase the level of congestion in the central part of the network.
- Airtime Fairness tests: in this case we repeated the iperf tests with and without the Airtime Fairness enabled.
- pingall test: in this experiment, a random subset of all the possible couples was taken and an iperf session was performed. This test should have been repeated for all the protocols, with and without Airtime fairness enabled, but was not possible to complete.

<i>Name</i>	<i>Description</i>
BABEL	Distance-vector LIII routing protocol
BATMAN Advanced v4	Source Routed LII routing protocol
BATMAN Advanced v5	Source Routed LII routing protocol (experimental)
BMX7	Source Routed LIII routing protocol
BMX7TUN	BMX7 with IP tunnel support
OLSRv1	Link state LIII routing protocol
OLSRv2	Link state LIII routing protocol
OLSRv2_MPR	OLSRv2 with MPR enabled

Tab. 1: The list of tested protocols

3 The Results

3.1 Ping Tests

Figure 3 reports the percentage of lost packets in the ping tests. It shows that OLSRv1 and BMX are the ones that choose paths that are more conservative, so they deliver all or almost all their pings. OLSRv2, BATMAN4 and, to a lower extent BABEL are the ones that, instead, tend to choose paths that are more lossy. BATMAN5 largely underperforms compared to the others. This is an example of a typical situation in the WBM. BATMAN5 does not have a stable release, and the developers brought to the WBM a testing version, to initially study its performance. During the tests, these outlier performances made it possible to spot the presence of previously unknown software bugs, which become visible only when tested on networks larger than a certain size. It was not possible to patch BATMAN5 before the end of the WBM and thus, from now on the results of BATMAN5 are omitted.

Figure 4 reports data about the distribution of the RTT measured with ping. What clearly emerges is that BMX is consistently using paths with larger delays, and that BABEL on three cases over 4 has a very large range of values. OLSRv2 performs worse than OLSRv1 in the majority of the cases, with delays comparable to BATMAN. Note that OLSRv2, BATMAN, and BABEL are advantaged in this comparison, since they have non-zero loss. It is reasonable to assume that the packets that could not be delivered, if they could reach the destination would probably have a large RTT.

Finally, figs. 5 to 8 report the whole set of RTT measured for each run. Note that the tests were done in sequence, so the network conditions may have varied from one test to the next one, note also the log scale on y axis. It is clear that there are two phenomena, one is the difference from one protocol to another that impacts the median value, the other is the presence of many outliers that influence the whiskers of fig. 4.

CDF diagrams for the same data (omitting BATMAN5, which adds too much noise to the figures) are presented in Figure 9.




Fig. 3: The percentage of lost packets per protocol

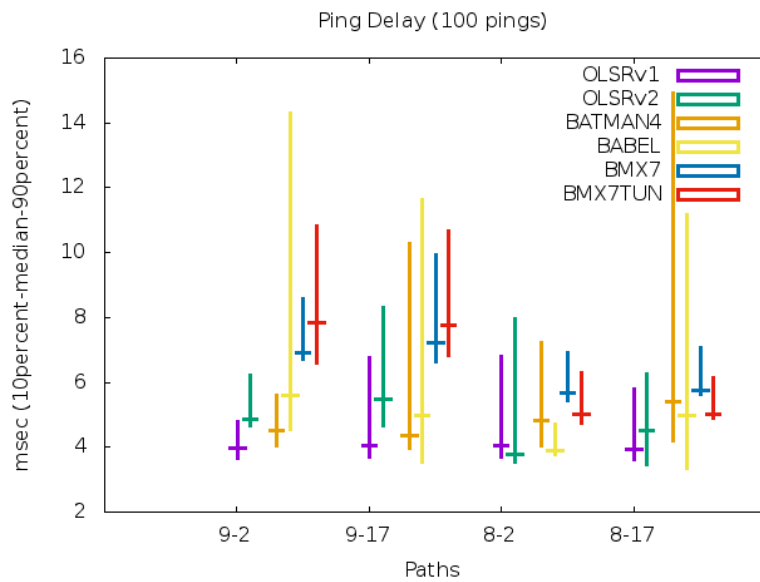


Fig. 4: The values of the 10th percentile, 90th percentile, and median value of the RTT for all the pings




Fig. 5: The values of the RTT per each ping, 9 to 17

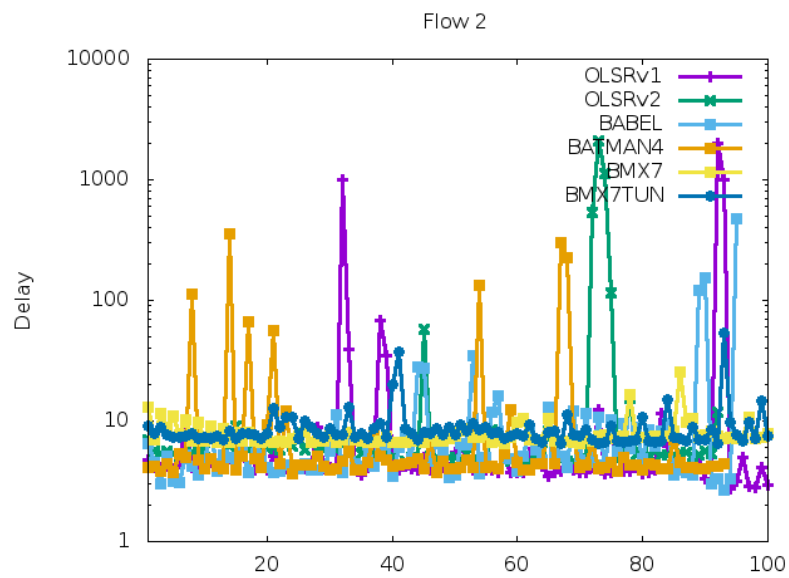


Fig. 6: The values of the RTT per each ping, 9 to 2

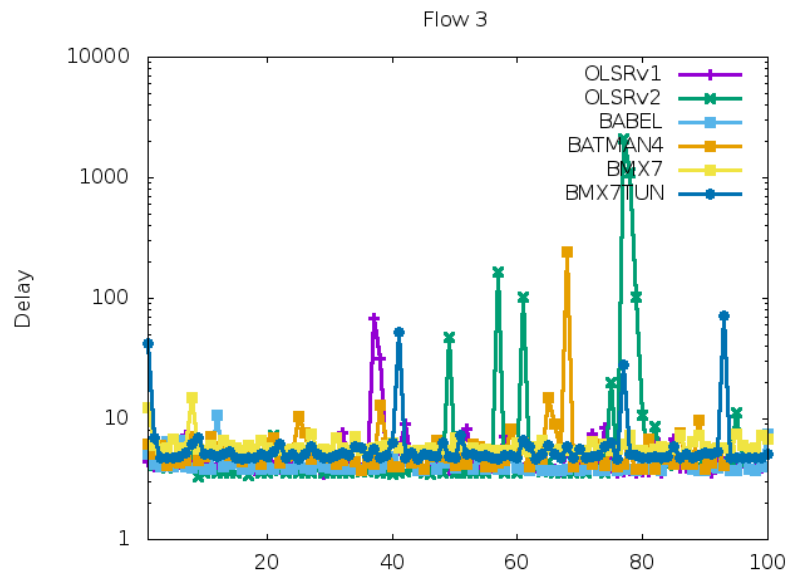


Fig. 7: The values of the RTT per each ping, 8 to 17




Fig. 8: The values of the RTT per each ping, 8 to 2

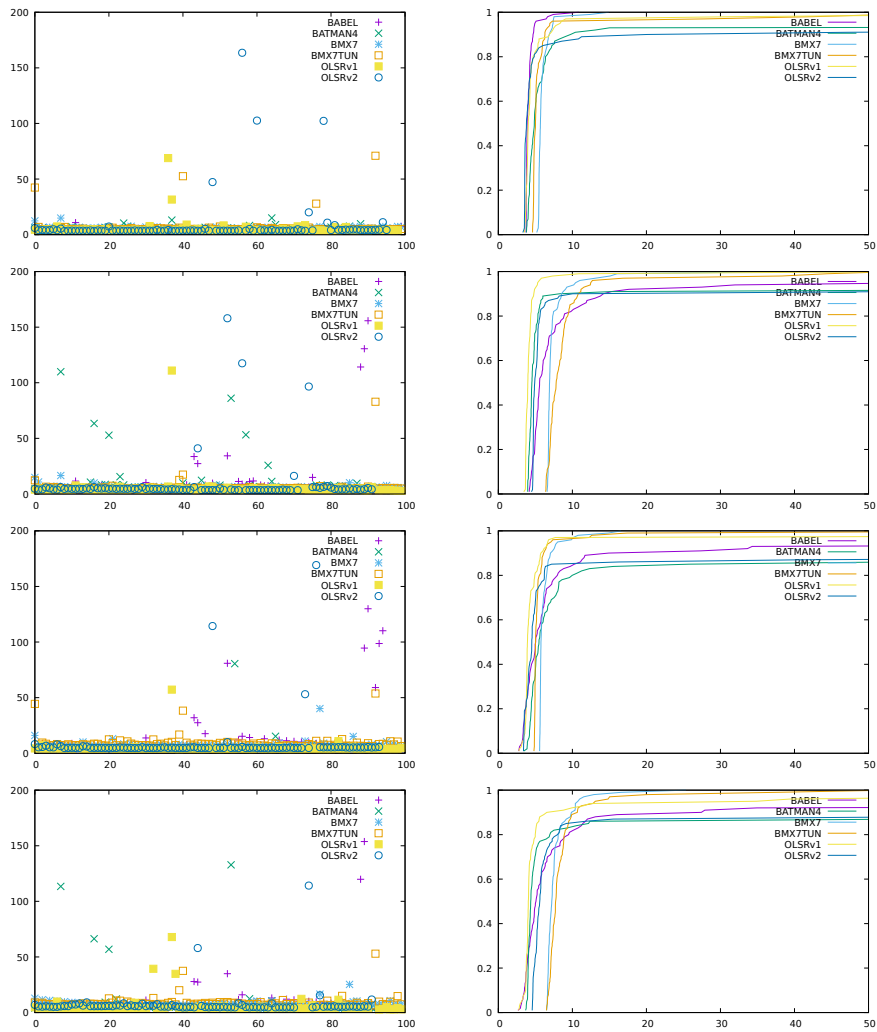


Fig. 9: Time-sequence and CDF plots of the ping data (links 8-17, 8-2, 9-17, 9-2).

3.2 iperf Test

Figures 10 to 12 report the performance of all the protocols when two parallel iperf sessions are run to measure the performance and zero, one or 4 sessions are added to increase background noise. The background sessions are fixed at 5Mbps and take place between couples (14,10), (4,15), (1,16), (12,11). Some observations that can be done are:

- There is a net performance degradation from one figure to the next one.
- With zero and one flow, in the path 8-17 (that according to fig. 2 is the shortest of the two) the protocol behave similarly, while in the path 9-2 there is a higher performance variation. In particular, BMX7 consistently performs better than the others (as the median value) but also shows the largest deviation. This may be the consequence of the combination of the zero loss and the stable delay that BMX7 is able to obtain (see figs. 3 and 4).
- with 5 flows, the relative difference increases and OLSRv2 seems to perform better than the other protocols averaging the results of both paths.

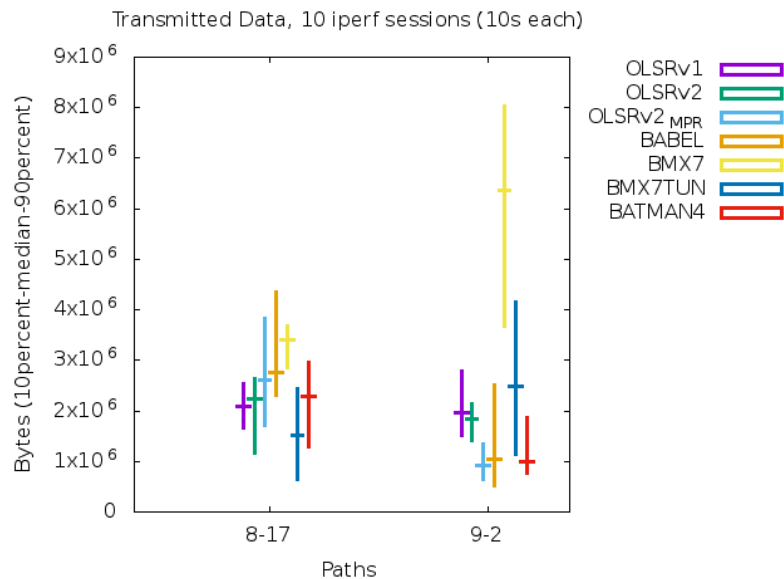


Fig. 10: iperf tests without background traffic.

Figure 13 reports the comparison of the performance of the iperf sessions, without background traffic with and without the Airtime fairness enabled. Each color in the figure corresponds to a protocol, the left value (with a dot at the median value) is the performance without the Airtime fairness enabled, the right value (with a dash at the median value) is the performance with the Airtime fairness enabled. It is hard to draw

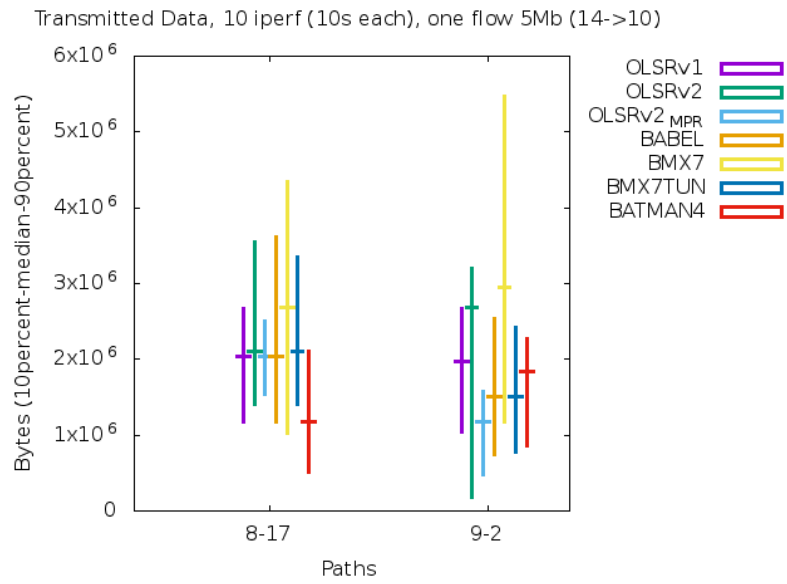


Fig. 11: iperf tests with one background traffic flow.

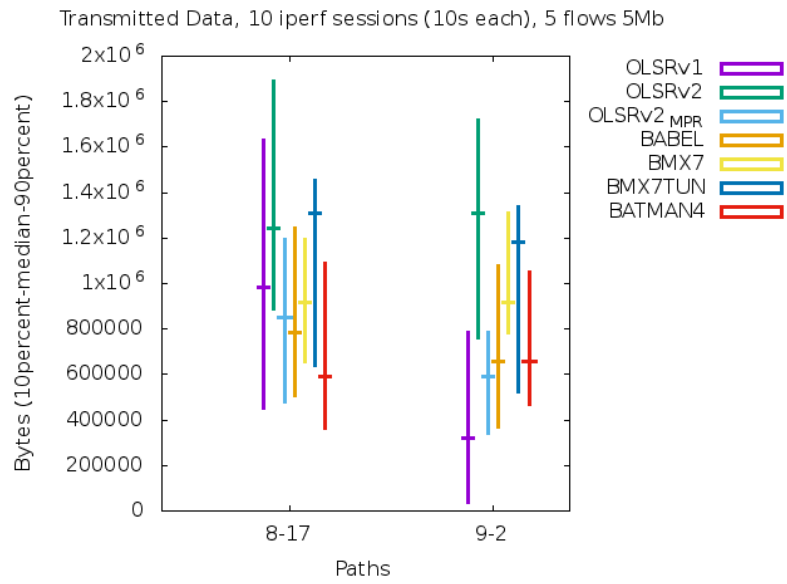


Fig. 12: iperf tests with four background traffic flows.

any conclusion on the importance of Airtime fairness, BMX7 still remains the protocol that guarantees the highest throughput, even if it is affected by the Airtime fairness in the opposite way in the two paths. Also the other protocols are affected both positively and negatively by Fairness, so it is unclear if the results are the effect of fairness, or of the changed conditions in the testbed from one run to the other. The only conclusion is that Airtime fairness does not influence the performance of mesh protocols in a critical way.

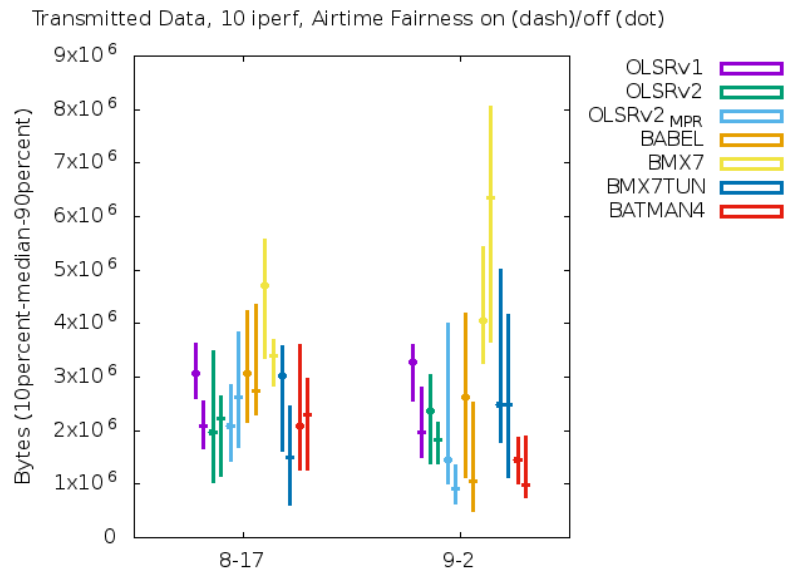


Fig. 13: iperf tests without background traffic and with/without airtime fairness.

4 Suggestion for the Future Battlemeshers

What follows is a set of recommendations we, as the people that spent most time on the testbed, want to give to the ones that in the next editions of the battlemesh will set-up the testbed.

Assign micro-tasks to other people

Setting up the testbed is full-time activity, so it is demanding on those that do it. Still, some tasks can be “outsourced” to others, tasks that are small, atomic and tedious, such as writing scripts to parse the output of commands used for testing (ping, iperf, traceroute...). When needed, instead of spending tens of minutes to write the script, it is useful to ask for help to others, this was done a couple of times and worked good. Ideally: write the task to an etherpad, ring a bell, and wait for someone to provide a script that does the job.

Use cables!

We choose 802.11s for the management network because, if it does not work, we can't blame any of the protocols under test. 802.11s works decently but, as any other solution is influenced by the state of the network. We believe that the best thing would be to buy a 300M reel of ethernet cable, and use that to manage the nodes. This would spare us of many failures in connecting, launching scripts, failing, restarting etc. . . . Powerline could be an alternative.

Do small tests

Large scale tests (all nodes against all nodes) are tempting, but since the network becomes complex it is hard to understand what is happening. We decided to start with small tests, where it is easier to understand what happens, while it happens.

Bring hardware for night-time tests

Probably the best moment to perform large-scale tests is during the night. But no one wants to leave his/her laptop at the BM nighttime. So it is important to remember to bring some other hardware (like a few raspberries) that can be used to perform automated tests during the night.

Fork the testbed

It happens that some of the devels need to use the testbed to fix bugs that happen during the experiments, and it is the added value of BM for them. On the other side, we can not freeze testing for long while they debug/improve their code. So a wise idea is to realize a small tbed (4-5 nodes) detached from the large one that devels can use to fix their bugs.

5 Conclusions

While there is a discussion in the BM community about changing the format of the event to something that is more conference-wise and less testbed-wise, our impression is that the testbed is important. Even if the results are incomplete, and scientifically not always sound, the whole process of setting up the testbed, running the protocols, debugging with the developers is an integral part of the conference and guarantees that the experts will participate to the conference, which makes it different from any other conference.