draft-ietf-bmwg-mlrsearch-02

IETF-113 Vienna BMWG Meeting

Authors: Vratko Polák, Maciek Konstantynowicz

Presented by: Vratko Polák

Draft Status

- In the middle of a rewrite.
- No changes to the algorithm logic (yet) since the previous draft version.
- Focusing on clearer terminology, less tight coupling between ideas, and making assumptions explicit.
- More reviews and comments are welcome.

Recap: IETF-109 slide

MLRsearch Sample Implementation

- A working implementation of MLRsearch is in Linux Foundation FD.io CSIT project.
  - Used for continuous measurements of NDR and PDR rates of:
    - FD.io VPP
    - DPDK L3fwd
    - DPDK Testpmd
  - Sample throughput results:
    - https://docs.fd.io/csit/rls2202/report/vpp_performance_tests/packet_throughput_graphs/index.html
  - General project info:
    - https://wiki.fd.io/view/CSIT
    - https://git.fd.io/csit/
- MLRsearch Python package (older version) published on PyPI:
  - https://pypi.org/project/MLRsearch/

Overview: Multiple Loss Ratio search (MLRsearch)

- MLRsearch is a packet throughput search algorithm suitable for deterministic systems.
- Non-zero target loss ratios.
  - More repeatable results than for zero loss ratio (Throughput).
  - Single search for all ratios at once.
- Uses shorter trial durations, so less trials at full duration are needed.
  - External search as a way to handle situation where longer duration trial result is not consistent with previous shorter duration trial results.
  - FRMOL-based initial trials, so not even shortest duration trials need to bisect the whole range.
- Smart load selection to save even more search time.

MLRsearch is compatible with RFC2544.

MLRsearch requirements

- Test equipment must be reliable:
  - The offered load must be close enough to the intended load.
  - Use smaller “max load” value to avoid loads where test equipment is not reliable.
  - Possibly do a “self test” to discover which loads are safe.
- Avoid very late packets (from previous trials):
  - RFC2544 already requires quite large delays.
- Search stop criteria:
  - Precision goal.
  - Timeout.
  - Optional early stop at min rate.

Real Example

- Table on the right shows data from a real test run in CSIT, using the default input values as described in the draft.
- The first column is the MLRsearch phase.
- The second is the trial measurement performed
  - Aggregate bidirectional offered load in mega (10^6) packets per second, and trial duration in seconds.
- Each of last four columns show one bound as updated after the measurement
  - Duration truncated to save space.
- Loss ratio is not shown, but invalid bounds are marked as N/A.
- **Bold** font signifies changed values.
- Blue font signifies results of the search.

Phase	Trial	NDR lower	NDR upper	PDR lower	PDR upper
init.	18.75 1.00	**N/A**	**18.75 1.**	**N/A**	**18.75 1.**
init.	6.14 1.00	N/A	**6.14** 1.	N/A	**6.14** 1.
init.	6.02 1.00	N/A	**6.02** 1.	**6.02 1.**	6.14 1.
int 1	5.56 1.00	**5.56 1.**	6.02 1.	6.02 1.	6.14 1.
int 1	5.78 1.00	**5.78** 1.	6.02 1.	6.02 1.	6.14 1.
int 1	5.90 1.00	5.78 1.	**5.90** 1.	6.02 1.	6.14 1.
int 2	5.84 5.48	5.78 1.	**5.84 5.**	6.02 1.	6.14 1.
int 2	5.78 5.48	**N/A**	**5.78** 5.	6.02 1.	6.14 1.
int 2	5.56 5.48	**5.56 5.**	5.78 5.	6.02 1.	6.14 1.
int 2	5.67 5.48	5.56 5.	**5.67** 5.	6.02 1.	6.14 1.
int 2	5.61 5.48	5.56 5.	**5.61** 5.	6.02 1.	6.14 1.
int 2	6.08 5.48	5.56 5.	5.61 5.	**6.08 5.**	6.14 1.
int 2	6.14 5.48	5.56 5.	5.61 5.	6.08 5.	**6.14 5.**
final	5.58 30.0	5.56 5.	**5.58 30**	6.08 5.	6.14 5.
final	5.56 30.0	**N/A**	**5.56** 30	6.08 5.	6.14 5.
final	5.45 30.0	N/A	**5.45** 30	6.08 5.	6.14 5.
final	5.03 30.0	**5.03 30**	5.45 30	6.08 5.	6.14 5.
final	5.23 30.0	**5.23** 30	5.45 30	6.08 5.	6.14 5.
final	5.34 30.0	5.23 30	**5.34** 30	6.08 5.	6.14 5.
final	5.28 30.0	5.23 30	**5.28** 30	6.08 5.	6.14 5.
final	5.26 30.0	**5.26** 30	5.28 30	6.08 5.	6.14 5.
final	6.11 30.0	5.26 30	5.28 30	6.08 5.	**6.11 30**
final	6.08 30.0	5.26 30	5.28 30	**6.08 30**	6.11 30

Fake Example

- Table on the right shows fake data.
- Linear scale, single doubling in external search.
- Loss ratio is not shown.
- **Bold** font signifies changed values.
- Blue font signifies results of the search.

Phase	Trial	NDRlo	NDRhi	PDRlo	PDRhi
init.	28 1.	**N/A**	**90 1.**	**N/A**	**28 1.**
init.	20 1.	N/A	**20** 1.	**20 1.**	28 1.
init.	16 1.	**16 1.**	20 1.	20 1.	28 1.
init.	24 1.	16 1.	20 1.	20 1.	**24** 1.
int 2	18 8.	16 1.	**18 8.**	20 1.	24 1.
int 2	16 8.	**16 8.**	18 8.	20 1.	24 1.
int 2	22 8.	16 8.	18 8.	**22 8.**	24 1.
int 2	24 8.	16 8.	18 8.	22 8.	**24 8.**
final	17 64	16 8.	**17 64**	22 8.	24 8.
final	16 64	**16 64**	17 64	22 8.	24 8.
final	23 64	16 64	17 64	22 8.	**23 64**
final	22 64	16 64	17 64	17 64	**22 64**
final	20 64	16 64	17 64	17 64	**20** 64
final	19 64	16 64	17 64	**19** 64	20 64

MLRsearch future improvements

- Obvious: Finish the current rewrite.
- Perform long duration trials for lower loss ratio before short duration trials for higher loss ratios.
  - I expect that will decrease search duration, maybe at the cost of more complicated load selection rules.
- Simplify load selection logic if possible.
- Allow different precision goals for different ratios (PDR can be more precise)?

THANK YOU !

draft-ietf-bmwg-mlrsearch-02

IETF-113 Vienna BMWG Meeting

Authors: Vratko Polák, Maciek Konstantynowicz

Presented by: Vratko Polák