draft-ietf-bmwg-mlrsearch-08

IETF-121 Dublin, BMWG Meeting

Authors: Vratko Polák, Maciek Konstantynowicz

MLRsearch Update

- draft-ietf-bmwg-mlrsearch-08 posted on 21st of October 2024.
- Changes from -07
  - Minor edits in MLRsearch Specification section – grammar, typos.
  - Major updates to Further Explanations section.
  - Fully updated / new Sample Search section with crispier references to RFC2544 and TST009 searches.
- BMWG next steps
  - Reviews.
  - Draft is ready for WG Last Call.

MLRsearch Update – Topics (content deltas from -07)

- Problems (recap)
- Specification
  - Existing Terms
  - Trial Terms
  - Goal Terms
  - Auxiliary Terms
  - Result Terms
  - Architecture
  - Compliance
- Further Explanations
  - Binary Search
  - Stopping Conditions and Precision
  - Loss Ratios and Loss Inversion
  - Exceed Ratio and Multiple Trials
  - Short Trials and Duration Selection
  - Generalized Throughput
- MLRsearch Logic and Example
  - Load Classification Logic
  - Conditional Throughput Logic
  - SUT Behaviors
  - Example Search
- Work Status as of -08
- References

Problems (recap)

- **Long Search Duration** => Let's shorten it by finding interesting region faster.
- **DUT in SUT** => Use performance spectrum (noiseless, noiseful) model to estimate DUT performance.
- **Repeatability and Comparability** => Identify consistent level of performance, find the noiseless end of the spectrum and reduce dependency on SUT noise.
- **Throughput with Non-Zero Loss** => Address a common practice in SW networking.
- **Inconsistent Trial Results** => Address the new phenomena resulting from repeated trials at the same load and multiple search goals.

Specification Update – Auxiliary Terms

Used for simplifying the language in discussion and explanation sections

- **Current and Final Quantities:** interim and final quantities computed and tracked by MLRsearch.
- **Trial Classification:** classification in relation to level of frame loss and trial duration
  - **High-Loss Trial:** a trial with Trial Loss Ratio larger than a Goal Loss Ratio value, for a given Search Goal.
  - **Low-Loss Trial:** a trial that is not a high-loss trial.
  - **Short Trial:** a trial with Trial Duration shorter than the Goal Final Trial Duration, for a given Search Goal.
  - **Full-Length Trial:** a trial that is not short trial.
  - **Long Trial:** a trial with Trial Duration longer than the Goal Final Trial Duration, for a given Search Goal.
- **Load Classification:**
  - **Upper Bound:** a Trial Load that is certain to fail to uphold all the requirements of the given Search Goal, mainly the Goal Loss Ratio in combination with the Goal Exceed Ratio.
  - **Lower Bound:** a Trial Load that is certain to uphold all the requirements of the given Search Goal.

Further Explanations 1/2

Further explanations of MLRsearch behavior in comparison with RFC 2544 bisection

- **Binary Search** - Recaps operation of [RFC2544] binary search tracking two tightest bounds, Upper Bound and Lower Bound, qualifies amount of time taken by search.
- **Stopping Conditions and Precision** - Clarifies criteria for stopping the Search.
- **Loss Ratios and Loss Inversion** - Explains strategies chosen to deal with inconsistent results from repeated measurements at the same load
  - **Single Goal and Hard Bounds** - Simplest case with single Search Goal and search always within “hard” bounds set by current lower bound and current higher bound.
  - **Multiple Goals and Loss Inversion** - Searching with multiple goals introduces the loss inversion phenomena as load that is upper bound for search goal-1 can be lower bound for search goal-2, and MLRsearch may perform trials at loads higher than upper bound for the search goal-1.
  - **Conservativeness and Relevant Bounds** - MLRsearch is conservative when dealing with Loss Inversion, the Upper Bound is considered real, and the Lower Bound is considered to be a fluke. Relevant Upper and Lower Bounds are tracking this per search goal and are analogous to tightest bounds in binary search.

Further Explanations 2/2

Further explanations of MLRsearch behavior in comparison with RFC 2544 bisection

- **Exceed Ratio and Multiple Trials** - Describes using multiple trials at the same load to discover noiseless and noiseful ends of the SUT performance spectrum. Exceed ratio facilitates results filtering by telling MLRsearch what ratio of trial duration seconds can have loss ratio larger than the goal loss ratio.
- **Short Trials and Duration Selection** – Clarifies usage and impact of short trials and makes recommendations to trial duration settings.
- **Generalized Throughput** – Provides further motivation and explanation of Conditional Throughput defined in MLRsearch.

RFC2544 Goal: Example Situation

- Example for RFC2544 Search Goal (unconditionally compliant):
  - Goal Final Trial Duration: **60 seconds**
  - Goal Duration Sum: **60 seconds**
  - Goal Loss Ratio: **0%**
  - Goal Exceed Ratio: **0%**
- Example situation during a Search:
  - Load: **1,000,000 frames per second**
  - One Trial measured (short, high-loss):
    - Trial Duration: **1 second** (less than Goal Final Trial Duration, so "short" trial)
    - Trial Loss Ratio: **0.1%** (more than Goal Loss Ratio, so "high-loss" trial)

RFC 2544 Goal: Load Classification Computation

1. Full-length high-loss sum:	**0s**	(no full-length trial)
2. Full-length low-loss sum:	**0s**	(no full-length trial)
3. Short high-loss sum:	**1s**	(the one trial is this)
4. Short low-loss sum:	**0s**	(no low-loss short trial)
5. Balancing sum:	**0s** = 0% * 0s	(exceed ratio of short low-loss sum)
6. Excess sum:	**1s** = 1s – 0s	(short high-loss sum minus balancing sum)
7. Positive excess sum:	**1s** = max(0, 1s)	(equal to excess sum as that is already positive)
8. Effective high-loss sum:	**1s** = 1s + 0s	(positive excess sum plus full-length high-loss sum)
9. Effective full sum:	**1s** = 1s + 0s	(effective high-loss sum plus full-length low-loss sum)
10. Effective whole sum:	**60s** = max(1s, 60s)	(max of effective whole sum and Goal Duration Sum)
11. Missing sum:	**59s** = 60s – 1s	(effective whole sum minus effective full sum)
12. Pessimistic high-loss sum:	**60s** = 1s + 59s	(effective high-loss sum plus missing sum)
13. Optimistic exceed ratio:	**1.67%** = 1s / 60s	(effective high-loss sum divided by effective whole sum)
14. Pessimistic exceed ratio:	**100%** = 60s / 60s	(pessimistic high-loss sum divided by effective whole sum)
15. Classification Result:	**Upper Bound**	(both exceed ratios above Goal Exceed Ratio)

TST009 Goal: Example Situation

- Example for TST009 Search Goal:
  - Goal Final Trial Duration: **60 seconds**
  - Goal Duration Sum: **120 seconds**
  - Goal Loss Ratio: **0%**
  - Goal Exceed Ratio: **50%**
- Example situation at the end of a Search:
  - Load: **1,000,000 frames per second**
  - Two Trials measured
  - First Trial (full-length, high-loss):
    - Trial Duration: **60 seconds** (not shorter than Goal Final Trial Duration, so "full-length" trial)
    - Trial Loss Ratio: **0.1%** (more than Goal Loss Ratio, so "high-loss" trial)
  - Second Trial (full-length, low-loss):
    - Trial Duration: **60 seconds** (still "full-length" trial)
    - Trial Loss Ratio: **0%** (not more than Goal Loss Ratio, so "low-loss" trial)

TST009 Goal: Conditional Throughput Computation

1. Full-length high-loss sum: **60s** (first Trial).
2. Full-length low-loss sum: **60s** (second Trial).
3. Full-length sum: **120s** = 60s + 60s (both trials).
4. Subceed ratio: **50%** = 100% - 50% (complement to Goal Exceed Ratio).
5. Remaining sum initially: **60s** = 50% * 120s (subceed ratio of full-length sum).
6. Current loss ratio initially: **100%**.
7. Sort Trial Results, 0% loss goes first.
8. Loop iteration 1, Trial duration 60s, loss 0%:
   - Remaining sum is **larger** than zero, not exiting the loop.
   - Set current loss ratio to this trial's Trial Loss Ratio which is **0%**.
   - Decrease the remaining sum by this trial's Trial Effective Duration.
   - New remaining sum is **0s** = 60s – 60s.
9. Loop iteration 2, Trial duration 60s, loss 0.1%:
   - Remaining sum is **not larger** than zero, exiting the loop.
10. Current loss ratio was most recently set to **0%**.
11. Current forwarding ratio is one minus the current loss ratio, so **100%**.
12. Conditional Throughput is the current forwarding ratio multiplied by the Load value.
13. Conditional Throughput is **one million frames per second**.

SUT behaviors

- With increasing Trial Duration, **probability** of next Trial being high-loss can:
  - Increase strongly:
    - Possible for both zero and non-zero Goal Loss Ratio.
    - Hints on resource leaks or large buffers.
    - Does not affect MLRsearch results.
    - Substantially increases search duration.
  - Increase mildly:
    - Common for zero Goal Loss Ratio.
    - Possible for non-zero Goal Loss Ratio, probable when infrequent loss bursts are large.
    - Does not affect MLRsearch results.
    - Somewhat increases search duration.
  - Stay constant:
    - Virtually impossible for zero Goal Loss Ratio.
    - Possible for non-zero Goal Loss Ratio, probable when infrequent loss bursts are small.
    - Does not affect MLRsearch results.
    - Leads to optimal search duration.
  - Decrease:
    - Not seen in practice, but possible if DUT has "cold start" each trial.
    - Impossible for zero Goal Loss Ratio, improbable even for non-zero Goal Loss Ratio, requires small loss bursts.
    - Reported MLRsearch results may be smaller than "true" result with no short trials, especially with zero exceed ratio.
    - Search duration is small with zero exceed ratio, larger than optimal otherwise.

MLRsearch Work Status

- Draft is ready for BMWG Last Call.

References

- MLRsearch is used in LFN FD.io CSIT
  - NDR Search goal
    - Goal Final Trial Duration: 1 seconds
    - Goal Duration Sum: 21 seconds
    - Goal Loss Ratio: 0%
    - Goal Exceed Ratio: 50%
  - PDR Search goal
    - Goal Final Trial Duration: 1 seconds
    - Goal Duration Sum: 21 seconds
    - Goal Loss Ratio: 0.5%
    - Goal Exceed Ratio: 50%
  - https://csit.fd.io
  - https://git.fd.io/csit/

THANK YOU !

draft-ietf-bmwg-mlrsearch-08

IETF-121 Dublin, BMWG Meeting

Authors: Vratko Polák, Maciek Konstantynowicz