Tunable Multi-Objective Tree Synthesis for Application-Layer Multicast

Kirill Karpov[1], Dmitry Kachan[1], Veronika Kirova[1], Ali Ghermezian[1], Daria Koshutina[2] and Eduard Siemens[1]

[1]*Anhalt University of Applied Sciences, Bernburger Str. 55, Koethen, Germany*

[2] *Odesa National Polytechnic University, Department of Information Systems, Shevchenko Ave. 1, Odesa, 65044, Ukraine*

{*Kirill.Karpov, Dmitry.Kachan, Veronika.Kirova, Ali.Ghermezian*}*@hs-anhalt.de, d.v.koshutina@op.edu.ua, Eduard.Siemens@hs-anhalt.de*

Keywords: Application Layer Multicast, RMDT, NSGA-II, Available Bandwidth, RTT, Minimum Spanning Tree, LAST, WAN.

Abstract: Application-layer multicast over wide-area networks must reconcile competing goals: paths should be short while links should offer high capacity and remain stable under variability. This paper presents a method that preserves multi-objective structure and exposes an explicit control of source-to-receiver latency. Edge preferences are derived from round-trip time (RTT) and available bandwidth (AvB) using NSGA-II and a seed tree is obtained by a minimum-spanning construction over the rank matrix. Path length is then constrained by a LAST-style adjustment evaluated in RTT, which guarantees that the distance to every receiver does not exceed a chosen factor α of the RTT shortest path. An equivalent additive budget L is supported for operational use. The approach is implemented with RMDT transport and evaluated on a five-region AWS testbed with active RTT/AvB probing and controlled cross-traffic. The NSGA-II–seeded trees consistently outperform single-metric baselines, and the latency constraint modifies only those segments that would otherwise incur excessive delay. The result is a monotone, single-parameter knob that lets operators satisfy latency requirements while retaining the performance benefits of multi-objective structure.

1 INTRODUCTION

Tree-based overlays are a common design for application-layer multicast, but their effectiveness depends on reconciling competing objectives. Minimizing path length is desirable, yet capacity, fairness, and robustness remain critical factors. The topology should reflect heterogeneous link qualities while remaining stable under variability and failures. In wide-area environments that are heterogeneous and time-varying, neither a fixed tree nor a single objective is sufficient. A tunable mechanism is therefore required that balances multi-objective structural preferences with the receiver path length, enabling selection along a continuum rather than commitment to a single operating point. This work adopts that view and develops a synthesis that preserves multi-objective structure while providing explicit, monotonic control of path length through a single parameter.

The proposed synthesis combines multi-objective ranking with path-length control. Edge preferences are derived from latency and available bandwidth using NSGA-II, and a tree is synthesized by applying a minimum-spanning procedure to the resulting rank matrix. Path length is then constrained through an α-controlled mechanism inspired by Light Approximate Shortest-path Trees. For every receiver v, the source-to-v latency in the resulting tree is guaranteed not to exceed α times the latency shortest-path distance. An equivalent additive formulation is exposed to operators: given a latency allowance L, the parameter is set as $\alpha = L/D_{\max}$, where $D_{\max}$ is the maximum latency shortest-path distance from the source. In this way, the multi-objective structure informed by both RTT and AvB is retained while path length is governed by a single, interpretable control.

The approach targets end-to-end multicast delivery with RMDT as the transport over the synthesized tree. Algorithmic design and complexity are discussed, and an implementation is provided. Evaluation is conducted on a global AWS testbed using instances deployed across five regions—US West (Oregon), EU (Frankfurt), EU (London), Asia Pacific (Singapore), a nd C anada (Central)—with multiple hosts per region, interface shaping to 100 Mbps, active probing of latency and available bandwidth,

and controlled cross-traffic. This environment provides reproducible, wide-area conditions for quantifying the achievable trade-off between multi-objective structure and path length.

The remainder of the paper is organized as follows: Section 2 surveys related work; Section 3 details the synthesis (NSGA-II ranks, MST seed, RTT-based LAST) and complexity; Section 4 outlines the AWS testbed and measurements; Section 5 reports results and decision-space visualizations; Section 6 concludes.

2 RELATED WORK

Multi-objective evolutionary optimization provides a natural machinery for combining heterogeneous link metrics. NSGA-II offers fast nondominated sorting, elitism, and a parameter-free diversity operator, and has become a standard baseline for bi-objective settings, including networking scenarios where latency and bandwidth conflict [1]. Its improved convergence and diversity make it suitable for constructing edge preferences from RTT and AvB.

Balancing path quality and global cost in spanning structures has been formalized through Light Approximate Shortest-path Trees (LAST): Khuller, Raghavachari, and Young proved that for $\alpha > 1$ a tree can simultaneously approximate shortest-path distances and MST weight, motivating post-processing of seed trees to enforce path-length guarantees [2].

Application-Layer Multicast (ALM) over WANs has been explored with reliable transport and measurement-informed construction. Prior work on RMDT analyzed multi-gigabit delivery and demonstrated benefits of ALM with active AvB estimation [3], arguing for transport/overlay co-design in the presence of cross-traffic. NSGA-II has further been applied to ALM to fuse RTT and AvB into a rank matrix used to build minimum-spanning trees, improving dissemination efficiency over single-metric designs [4]. The present work complements these results by adding explicit operator control over source-to-receiver path length via a LAST-style relaxation.

Beyond classical ALM, recent studies examine metric- and constraint-aware overlays: Cloudcast (NSDI'24) constructs cost-aware multicast trees using an egress-budget knob [5]; [6] builds low-delay, high-stability trees via a fan-out-to-delay heuristic; and [7] uses a discrete Artificial Fish Swarm Algorithm to jointly optimize multiple ALM trees under degree constraints. In SDN, [8] adapts multicast construction to device capacity. Complementing these, hierarchical clustering-based multicast [9] and fairness-oriented routing for distributed interactive applications [10] highlight the role of multi-criteria routing and latency/fairness requirements. In contrast to these approaches, we fuse RTT and AvB via NSGA-II and then apply a LAST-style, operator-facing RTT bound, yielding a single-parameter mechanism that preserves multi-objective structure while enforcing latency constraints.

3 METHODOLOGY

Let $G = (V,E)$ be an undirected overlay with two edge metrics, round–trip time $\mathrm{RTT}(e)$ and available bandwidth $\mathrm{AvB}(e)$, and let $r \in V$ be the source. The goal is to synthesize a multicast tree T that reflects a multi–objective structure derived from both metrics while enforcing a controllable bound on source–to–receiver path length measured in RTT.

Multi–objective preferences are first computed per edge by applying NSGA-II to the pairs $\big(\mathrm{RTT}(e), \mathrm{AvB}(e)\big)$. The outcome is a symmetric *rank matrix* R that orders edges by Pareto dominance (low latency, high bandwidth). A minimum–spanning procedure is then applied to R to obtain a seed tree

$$T_0 \;=\; \mathrm{MST}\big(G; \mathrm{weight} = R\big),$$

which encodes the structural balance between the two objectives without committing to either one alone.

Path length is controlled next in the latency metric. Single–source shortest–path distances $d_{\mathrm{SPT}}^{\mathrm{RTT}}(\cdot)$ from r are computed once on G with RTT as edge weight. The seed T_0 is processed by an α–constrained Light Approximate Shortest–path Tree (LAST) pass evaluated in RTT: vertices are examined in depth–first order, and whenever the current latency distance in the working tree exceeds $\alpha d_{\mathrm{SPT}}^{\mathrm{RTT}}(v)$, the RTT–shortest path from r to v is inserted and the unique cycle edges are removed. The result is which preserves as much of the NSGA–derived structure as possible while meeting the latency constraint for every receiver.

Operators may specify the constraint either multiplicatively, via $\alpha \geq 1$, or additively, via an allowance L (milliseconds). The two forms are related by

$$d_{T_\alpha}^{\mathrm{RTT}}(v) \leq \alpha \cdot d_{\mathrm{SPT}}^{\mathrm{RTT}}(v) \Leftrightarrow d_{T_\alpha}^{\mathrm{RTT}}(v) \leq d_{\mathrm{SPT}}^{\mathrm{RTT}}(v) + L,$$
$$\alpha \;=\; \frac{L}{D_{\max}}, \tag{1}$$

where $D_{\max} = \max_{u \in V} d_{\mathrm{SPT}}^{\mathrm{RTT}}(u)$. The distance guarantee holds by construction because all checks and insertions are performed in RTT. Since the seed is an MST over NSGA–II ranks rather than over latency, classical bounds on total latency cost relative to an

RTT-MST are not asserted; aggregate costs are reported empirically.

Computationally, NSGA-II ranking over all pairs induces $O(n^2 \log n)$ preprocessing (implementation–dependent). The MST over R runs in $O(m \log n)$, the single–source RTT shortest paths run in $O(m \log n)$, and the LAST pass is linear in the number of inserted path edges and is near $O(m)$ on sparse overlays.

4 EXPERIMENTAL SETUP

Experiments are conducted on a global Amazon AWS testbed. In each of five regions - US West (Oregon), EU (Frankfurt), EU (London), Asia Pacific (Singapore), and Canada (Central) three `c5.xlarge` instances are deployed, yielding a geographically diverse overlay with both intra- and inter-continental paths. All hosts run Ubuntu 24.04 LTS on 4 vCPUs and 8 GB RAM with a 10 Gbps virtual NIC; to remove host- and NIC-side bottlenecks and to control link utilization, interface rates are shaped to 100 Mbps using Linux `tc`. Regions are color-coded consistently across figures: Oregon orange, Frankfurt blue, London red, Singapore cyan, and Canada green. Nodes are labeled by a region initial and an index (e.g., *l1*, *c2*, *o3*), where *o*, *l*, *c*, *f*, and *s* denote Oregon, London, Canada, Frankfurt, and Singapore, respectively; the numeral is the VM index within the region. The host geography is shown in Fig. 1.

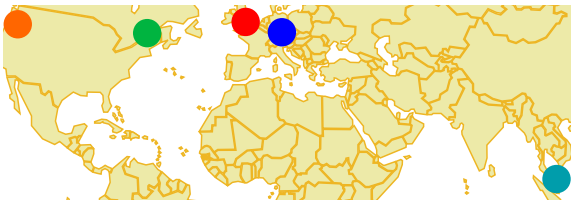

Figure 1: Location of AWS hosts used in the experiments.

Latency (RTT) and available bandwidth (AvB) are measured actively between all host pairs. RTT is obtained via ICMP probing; AvB is estimated using the `Kite2` active probing tool suitable for high-speed links. The resulting RTT matrix used by the synthesis is shown as a heatmap in Fig. 2.

To study behavior under interference, controlled cross-traffic is generated with `iperf`, producing the load pattern in Fig. 3; this configuration induces cases where bandwidth-aware selections diverge from delay-only choices and is used to evaluate multi-objective synthesis and path-length control.

	l1	l2	l3	f1	f2	f3	c1	c2	c3	o1	o2	o3	s1	s2	s3
l1	0.0	0.1	0.1	13.4	14.3	14.2	89.1	85.6	88.8	133.6	137.4	133.7	140.8	140.2	140.4
l2	0.1	0.0	0.1	13.3	15.3	13.0	85.6	86.9	87.3	138.6	133.7	137.8	185.2	181.2	184.7
l3	0.1	0.1	0.0	13.1	13.7	13.6	88.8	88.8	86.8	137.1	139.0	139.9	180.7	181.2	171.9
f1	13.4	13.3	13.1	0.0	0.1	0.1	100.4	100.4	98.9	158.5	161.1	164.3	177.7	169.6	173.7
f2	14.3	15.3	13.7	0.1	0.0	0.1	99.3	99.7	99.2	158.9	162.6	162.7	173.3	173.4	173.7
f3	14.2	13.0	13.6	0.1	0.1	0.0	100.4	99.3	98.9	163.2	158.3	163.9	173.6	174.0	173.0
c1	89.1	85.6	88.8	100.4	99.3	100.4	0.0	0.1	0.1	66.5	66.6	65.1	219.8	219.9	219.8
c2	85.6	86.9	88.8	100.4	99.7	99.4	0.1	0.0	0.1	65.6	66.6	67.1	219.9	219.7	219.7
c3	88.8	87.3	86.8	98.9	99.2	98.9	0.1	0.1	0.0	67.4	66.6	66.7	219.9	219.6	219.6
o1	133.6	138.6	137.1	158.5	159.1	163.2	66.5	65.6	67.4	0.0	0.3	0.3	163.2	162.5	162.6
o2	137.4	133.7	139.0	161.2	162.6	158.2	66.6	66.6	66.6	0.3	0.0	0.3	163.1	162.9	162.8
o3	133.7	137.8	139.9	164.3	162.7	163.8	65.1	67.1	66.6	0.3	0.3	0.0	163.5	163.1	162.6
s1	140.8	185.2	180.7	177.7	173.3	173.6	219.7	219.9	219.9	163.2	163.1	163.5	0.0	0.1	0.1
s2	140.2	181.2	181.2	169.6	173.4	174.0	219.9	219.7	219.6	162.5	162.9	163.1	0.2	0.0	0.1
s3	140.4	184.0	171.9	173.7	173.7	173.0	219.8	219.7	219.6	162.6	162.8	162.6	0.1	0.1	0.0

Figure 2: RTT matrix (ms) between all hosts; rows denote sources and columns denote destinations. Cell styling marks spanning-tree selections: bold red entries indicate edges chosen by the RTT-based MST, red-bordered entries indicate the AvB-based MST, and green-shaded entries indicate the multi-objective MST built from NSGA-II ranks.

	l1	l2	l3	f1	f2	f3	c1	c2	c3	o1	o2	o3	s1	s2	s3
l1	100.0	100.0	100.0	86.7	86.7	86.6	94.5	94.9	94.6	91.9	91.7	91.8	86.7	86.7	86.7
l2	100.0	99.9	100.0	86.6	94.7	86.7	100.1	94.8	94.7	91.7	91.9	91.7	100.0	100.1	100.0
l3	99.9	100.0	100.1	86.7	86.7	86.7	94.6	94.7	94.7	91.7	91.5	91.5	100.0	100.0	99.9
f1	86.6	86.7	86.6	100.1	100.0	100.0	88.8	88.8	89.1	90.0	90.1	90.1	89.3	89.4	89.4
f2	86.7	94.7	86.7	100.0	99.9	99.9	89.0	89.0	89.0	90.2	90.1	90.0	89.7	89.6	89.4
f3	86.7	86.8	86.7	100.0	100.1	100.0	89.6	88.9	89.6	90.1	90.2	90.3	89.4	89.5	89.4
c1	94.6	94.8	94.7	88.8	88.8	89.0	100.0	100.0	100.0	86.6	86.6	86.5	93.9	94.0	93.9
c2	94.8	94.8	94.6	89.0	88.9	88.9	100.0	99.9	100.0	86.6	94.7	86.7	94.0	93.9	94.0
c3	94.6	94.7	94.7	89.6	88.9	89.6	100.0	100.1	100.0	86.6	86.6	86.7	93.9	93.9	94.1
o1	91.8	91.6	91.7	90.1	90.1	90.1	86.6	86.6	86.7	100.0	100.0	100.0	90.4	90.3	90.1
o2	91.7	91.9	91.6	90.1	90.1	90.1	86.7	94.7	86.7	100.0	100.0	100.0	90.2	90.1	90.4
o3	91.9	91.7	91.4	90.1	90.1	90.1	86.7	86.8	86.6	100.0	100.0	99.9	90.0	90.2	90.0
s1	86.6	100.0	100.0	89.2	89.7	89.5	93.9	93.9	94.0	90.4	90.2	90.1	100.0	100.0	100.1
s2	86.6	100.0	100.1	89.4	89.5	89.5	94.1	94.0	94.0	90.4	90.2	90.1	100.1	100.0	100.0
s3	86.6	100.0	100.1	89.6	89.4	89.6	94.0	94.0	94.0	90.1	90.4	90.1	100.0	100.0	100.0

Figure 3: Estimated AvB matrix in Mbps; rows denote source, columns denote destination hosts. Cell styling marks MST selections: red entries indicate edges chosen by the RTT-based MST, red-bordered entries indicate the AvB-based MST, and green entries indicate the multi-objective MST built from NSGA-II ranks.

5 EXPERIMENTAL RESULTS

The efficacy of each routing strategy is evident when comparing the multicast trees produced under four different optimization criteria: minimizing RTT, maximizing AvB, multi-objective NSGA-II ranking, and the NSGA-II variant with a latency constraint. A color-coded latency matrix (see Figure 4) annotates the links selected by each approach, using bold red text to denote edges in the RTT-based MST, red-bordered cells for edges in the AvB-based MST, green shaded cells for edges chosen by the NSGA-II MST, and orange shaded cells for those in the latency-constrained (LAST) tree. This visualization allows a direct interpretation of how each strategy influences path selection across the network.

The RTT-optimized MST favors the shortest paths between regions, as expected. The tree built on RTT immediately links distant nodes via the minimum-latency route – notably connecting London directly to Singapore – to minimize end-to-end delay. In Figure 5, the London–Singapore link indicates that the RTT-based tree includes this direct 140ms path. This strategy yields a low-delay distribution structure but can overlook link capacity; many of its

	l1	l2	l3	f1	f2	f3	c1	c2	c3	o1	o2	o3	s1	s2	s3
l1	0.0	0.1	0.1	13.4	14.3	14.2	89.1	85.6	88.8	133.6	137.4	133.7	140.8	140.2	140.4
l2	0.1	0.0	0.1	13.3	15.3	13.0	85.6	86.9	87.3	138.6	133.7	137.8	185.2	181.2	184.7
l3	0.1	0.1	0.0	13.1	13.7	13.6	88.8	88.8	86.8	137.1	139.0	139.9	180.7	181.2	171.9
f1	13.4	13.3	13.1	0.0	0.1	0.1	100.4	100.4	98.9	158.5	161.1	164.3	177.7	169.6	173.7
f2	14.3	15.3	13.7	0.1	0.0	0.1	99.3	99.7	99.2	158.9	162.6	162.7	173.3	173.4	173.7
f3	14.2	13.0	13.6	0.1	0.1	0.0	100.4	99.3	98.9	163.2	158.3	163.9	173.6	174.0	173.0
c1	89.1	85.6	88.8	100.4	99.3	100.4	0.0	0.1	0.1	66.5	66.6	65.1	219.8	219.9	219.8
c2	85.6	86.9	88.8	100.4	99.7	99.4	0.1	0.0	0.1	65.6	66.6	67.1	219.9	219.7	219.7
c3	88.8	87.3	86.8	98.9	99.2	98.9	0.1	0.1	0.0	67.4	66.6	66.7	219.9	219.6	219.6
o1	133.6	138.6	137.1	158.5	159.1	163.2	66.5	65.6	67.4	0.0	0.3	0.3	163.2	162.5	162.6
o2	137.4	133.7	139.0	161.2	162.6	158.2	66.6	66.6	66.6	0.3	0.0	0.3	163.1	162.9	162.8
o3	133.7	137.8	139.9	164.3	162.7	163.8	65.1	67.1	66.6	0.3	0.3	0.0	163.5	163.1	162.6
s1	140.8	185.2	180.7	177.7	173.3	173.6	219.7	219.9	219.9	163.2	163.1	163.5	0.0	0.1	0.1
s2	140.2	181.2	181.2	169.6	173.4	174.0	219.9	219.7	219.6	162.5	162.9	163.1	0.2	0.0	0.1
s3	140.4	184.0	171.9	173.7	173.7	173.0	219.8	219.7	219.6	162.6	162.8	162.6	0.1	0.1	0.0

Figure 4: RTT matrix in ms; rows denote sources and column denote destination hosts. Cell styling marks MST selections: red entries indicate edges chosen by the RTT-based MST, red-bordered entries indicate the AvB-based MST, green entries indicate the multi-objective MST built from NSGA-II ranks, and orange entries indicate the NSGA-II–seeded tree after LAST-based relaxation with latency allowance L=150 msec).

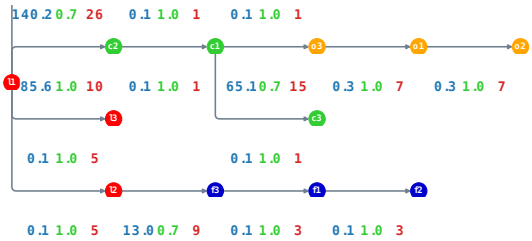

Figure 5: MST computed on the latency metric (RTT). Edge annotations show RTT in ms (blue), available bandwidth (normalized, green), and multi-objective rank (red).

chosen edges are short-haul connections (e.g., London–Frankfurt at ∼13ms) that prioritize latency over bandwidth. In contrast, the AvB-optimized MST produces a very different topology (Figure 3), selecting links that maximize available throughput even at the expense of longer physical routes. The red-bordered annotations in the matrix highlight how the AvB-based tree avoids certain low-latency links. In the critical London–Singapore segment, the AvB strategy forgoes the direct London–Singapore path (despite its latency advantage) in favor of an alternate route with higher capacity – for instance, routing via the Oregon region. This results in a larger RTT (the trans-Pacific path London→Oregon→Singapore is notably longer than the direct path), underscoring the conflict between the RTT and AvB metrics in the l→s region. In general, the AvB-MST connects distant sites through high-bandwidth backbone links (e.g., inter-continental links through North America, marked with red borders in the matrix) rather than the shortest geographic paths. The divergence between the RTT- and AvB-optimized trees confirms that no single metric is universally optimal: the RTT-based tree achieves low delay to all nodes, whereas the AvB-based tree offers higher per-link capacity but can significantly inflate latency on certain paths (in this case, to Singapore).

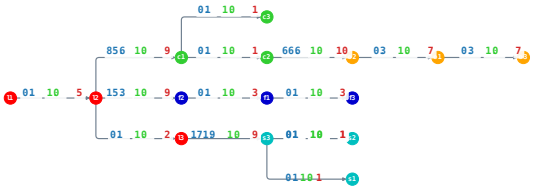

Figure 6: MST computed on the multi-objective NSGA-II rank metric. Edge annotations show RTT in ms (blue), available bandwidth (normalized, green), and NSGA-II rank (red).

The multi-objective NSGA-II approach yields a spanning tree that strikes a balance between these extremes. As shown by the green-highlighted cells in the matrix and the network diagram (Figure 6), the NSGA-II optimized MST tends to choose alternative routes in regions where RTT and AvB trade off most sharply. Notably, the NSGA-II tree does *not* simply mirror the RTT-minimal solution for the London–Singapore connection. Instead, it often routes traffic to Singapore through an intermediate hub (the Oregon node) which offers substantially higher available bandwidth, incurring only a moderate RTT increase (the l→s one-way delay rises to ∼163ms via Oregon, instead of 140ms direct). This compromise is evident from the matrix: the direct l–s link that was bold red in the pure RTT tree is no longer selected by NSGA-II; instead, a different cell along the l→o→s path is shaded green, indicating the NSGA-II MST's preference for that higher-bandwidth route. In effect, the NSGA-II algorithm identifies links that provide a favorable trade-off – generally those with reasonably low RTT *and* high AvB – and the resulting tree avoids the worst cases of either metric. Indeed, when mapping all inter-host links on an RTT-vs-AvB plot, the NSGA-II's chosen edges cluster in the top-left quadrant (low latency, high bandwidth), confirming its balanced optimization.

Figure 7 plots all candidate inter-host links in the RTT–AvB plane. NSGA-II selections concentrate in the upper-left portion (low RTT, high AvB); for London→Singapore the chosen edge lies in the long-range band at ≈160 ms but close to the high-AvB plateau, consistent with the trade-off observed in Figure 6. Performance measurements reflect this balance: the NSGA-II derived tree sustained a significantly higher average data rate than either single-metric tree, which achieved about 32Mbps (RTT-based) and 20 Mbps (AvB-based) per edge. The NSGA-II MST's ability to find an alternate route for the Singapore region while keeping most links short and uncongested illustrates its effectiveness in jointly minimizing delay and maximizing throughput across the network.

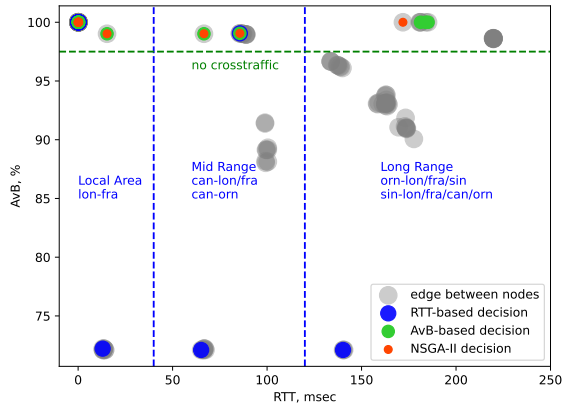

Figure 7: Pre-LAST decision space for candidate edges.

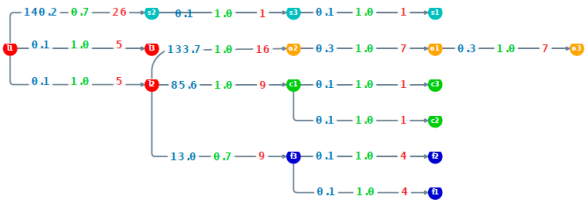

Figure 8: Tree synthesized by α-constrained relaxation of the NSGA-II–rank MST toward RTT shortest paths. Edge annotations show RTT in ms (blue), AvB (normalized, green), and NSGA-II rank (red).

Finally, imposing a latency constraint on the multi-objective solution (the LAST approach) alters the tree in precisely those scenarios where NSGA-II was willing to tolerate longer delays. By introducing an RTT threshold of $L = 150$msec, we effectively forbid extremely long-delay links from appearing in the spanning tree, while still leveraging the NSGA-II ranking for multi-factor optimization. The orange-shaded entries in the annotated matrix correspond to the edges of the resulting LAST-constrained tree (Figure8). In the London–Singapore case, the effect is to re-introduce a latency-aware path: the direct London–Singapore link (140ms) reappears in the final tree, highlighted in orange, replacing the two-hop high-bandwidth route that exceeded the 150ms RTT limit. In other words, the LAST strategy converges toward the RTT-optimal choice for that particular leg, ensuring the Singapore hosts are reached without violating the latency budget. More generally, the LAST-based relaxation yields a topology very similar to the original NSGA-II MST, but with any overly long links pruned out in favor of slightly slower (in bandwidth) but faster (in RTT) alternatives. The Oregon–Singapore link (163ms) is one such edge that the unconstrained NSGA-II tree had selected for its bandwidth; under the 150ms cap, that edge is replaced by the direct London–Singapore connection. This adjustment has minimal impact on overall throughput – the revised tree still delivers high bandwidth on most segments – but it guarantees that no path in the multicast tree suffers from extreme latency. The result is a latency-aware multi-objective tree that preserves the bulk of NSGA-II's performance gains while adhering to practical RTT limits.

Figure 9 shows the same decision space with the red vertical line marking the path-length limit $L = 150$ ms (converted to α). All red markers (LAST+NSGA-II selections) fall to the left of this line; in particular, the Oregon detour at ≈163 msec is discarded and the direct London→Singapore edge at 140 ms becomes selected. The remaining choices stay near the high-AvB band, indicating that the latency constraint tightens only the longest paths while preserving the multi-objective structure.

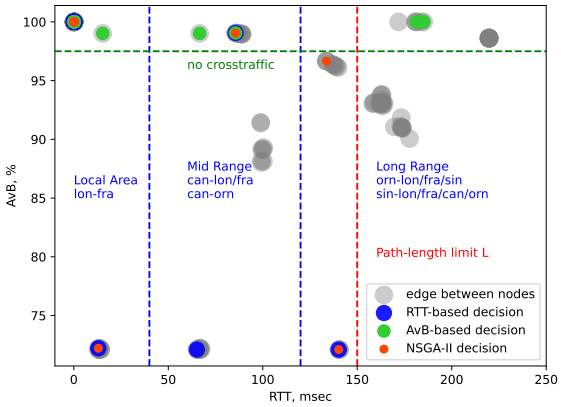

Figure 9: Post-LAST decision space for candidate ALM edges. RTT metric limited by 150 msec.

The NSGA-II optimized strategy chooses well-balanced routes that outperform single-metric designs, and the LAST refinement further ensures that even the longest paths (especially the challenging l→s route) remain within acceptable delay, all without drastically compromising the bandwidth benefits of the multi-objective solution.

6 CONCLUSIONS

A tunable synthesis for application-layer multicast was presented that combines multi-objective edge ranking with explicit path-length control. Edge preferences derived from RTT and available bandwidth via NSGA-II seed a minimum-spanning tree that captures the structural trade-off between latency and capacity; a subsequent LAST-style relaxation in the

RTT metric enforces per-receiver latency within a chosen bound, expressed either multiplicatively (α) or as an additive allowance L. Experiments on a five-region AWS testbed showed that the NSGA-II seeded tree consistently yields better-balanced routes than single-metric designs, improving dissemination efficiency while avoiding the worst cases of either RTT-only or AvB-only selection. Introducing a latency allowance (e.g., $L = 150$ ms) adjusts only those segments where the multi-objective seed tolerates excessive delay, as illustrated by the London→Singapore path: the unconstrained NSGA-II tree prefers a higher-bandwidth detour (via Oregon, ≈ 163 ms), whereas the constrained variant reverts to the direct link (140 ms) with negligible impact on overall throughput.

Operationally, the approach provides a single, interpretable control that aligns with service objectives: an operator may state a latency requirement (e.g., "no branch worse than 150 ms over shortest paths") and obtain a tree that preserves the multi-objective structure while satisfying that constraint. Complexity remains near that of standard graph primitives (MST and SPT), and the method integrates cleanly with reliable ALM transport such as RMDT. The results indicate that even when a globally balanced solution exists, local path exceptions can be corrected by a simple limit without sacrificing the advantages of multi-objective optimization.

ACKNOWLEDGEMENTS

This work was supported by the European Regional Development Fund (ERDF/EFRE) and the State of Saxony-Anhalt within the programme *Sachsen-Anhalt WISSENSCHAFT Forschung und Innovation (EFRE) 2021–2027*, project AI-RMDT (grant no. ZS/2023/12/182323). We acknowledge support by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) and the Open Access Publishing Fund of Anhalt University of Applied Sciences.

REFERENCES

[1] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," IEEE transactions on evolutionary computation, vol. 6, no. 2, pp. 182–197, 2002.

[2] S. Khuller, B. Raghavachari, and N. Young, "Balancing minimum spanning trees and shortest-path trees," Algorithmica, vol. 14, no. 4, pp. 305–321, 1995.

[3] K. Karpov, V. Kirova, N. Mareev, D. Syzov, and E. Siemens, "Data transmission performance enhancement in multi-gigabit wide area networks," in Advances in Information and Communication Technology and Systems. Cham: Springer International Publishing, 2021, pp. 161–190.

[4] K. Karpov, D. Kachan, and E. Siemens, "Enhancing application layer multicast using multi-objective optimization techniques," in Information and Communication Technologies and Sustainable Development. Cham: Springer Nature Switzerland, 2023, pp. 42–53.

[5] S. Wooders, S. Liu, P. Jain, X. Mo, J. E. Gonzalez, V. Liu, and I. Stoica, "Cloudcast: High-throughput cost-aware overlay multicast in the cloud," in 21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24), 2024, pp. 281–296.

[6] D. Li, Z. Wang, Y. Wei, J. Yao, Y. Tan, Q. Yang, Z. Wang, and X. Cao, "Generation of low-delay and high-stability multicast tree." Computers, Materials & Continua, vol. 76, no. 1, 2023.

[7] Y. Li, N. Wang, W. Zhang, Q. Liu, and F. Liu, "Discrete artificial fish swarm algorithm-based one-off optimization method for multiple co-existing application layer multicast routing trees," Electronics, vol. 13, no. 5, 2024.

[8] G. Bijur, M. Ramakrishna, and K. A. Kotegar, "Multicast tree construction algorithm for dynamic traffic on software defined networks," Scientific Reports, vol. 11, no. 1, p. 23084, 2021.

[9] K. Somasundaram and P. R. Kanna, "Scalable hierarchical balanced clustering-based routing with multipath authentication for secured data transmission in large-scale multicast group communications," Expert Systems with Applications, vol. 286, p. 128149, 2025.

[10] I. F. Ferreira, M. L. M. Peixoto, and G. B. Figueiredo, "Fairness-oriented multicast routing for distributed interactive applications," Computer Communications, vol. 219, pp. 229–242, 2024.