Core-Selection Algorithms in Multicast Routing

Ayşe Karaman and Hossam Hassanein

School of Computing
Queen's University
Kingston, ON K7L 3N6, Canada
{karaman, hossam}@cs.queensu.ca

Abstract

The Core-based approach is inevitable in multicast routing protocols as it provides efficient management of multicast path in changing group memberships, and scalability and performance. In this paper, we present a comprehensive analysis of this approach with the emphasis on core selection for the first time in literature. We first examine the evolution of multicast routing protocols into the core-based architecture and the motivation for the approach. Then we review the core-selection algorithms in the literature for their algorithmic structure and performance. Our study involves an extensive computational and message complexity analysis of each algorithm, and a classification for their deployment characteristics and algorithmic complexities. To the best of our knowledge ours is the first paper providing such extensive comparative analysis of core-based multicast routing protocols.

1. Introduction

Multicasting is the simultaneous delivery of data stream to multiple receivers in the domain. Its applications are numerous in group communications including numerous forms of audiovisual conferencing (discussion forums, live auctions, negotiation systems, white-board applications, online games, chat-rooms) and broadcasting (online radios, e-learning systems), mass delivery in digital marketing (software upgrade distribution, news delivery, stock quotes), replicated database applications in various forms of cross-query systems (airline systems, banking applications, cross-catalog search), offline virtual communities (newsgroups, mailing lists), as well as the trivial resource discovery feature of Internet routers. Group communications on Internet are increasingly more pervasive, and are inevitable parallel to wider acceptance of ecommerce applications.

The objective of multicast routing, in broad terms, is the efficient delivery of the stream to the recipients. Multicast routing, as opposed to unicast routing, attempts to duplicate copies of the packets only at branching nodes, nodes where there are more than one outgoing paths to multiple destinations. Group communications have varying demands on network resources and transmission capacities. Some applications require reliability (B2B transactions, mailing lists) while others demand for speed (B2C services, online broadcasting) or cost efficiency of

transmission (delivery sy

routing protocol, on the other hand, is the entire architecture that includes the multicast tree building algorithm as well as the maintenance of the tree structure in the dynamisms of the multicast environment, packet forwarding, failure discovery and recovery, and other features depending on the protocol specifics. The main architectural component of a multicast routing protocol is the development of the multicast tree not only because it is the structure that directly affects the performance of the protocol, but also is the most demanding one for the development and maintenance. nother major concern in multicast protocol design is the feasibility of the deployment while maintaining high performance in the presence of cost of control overhead to operate and maintain the construct to support that performance. These two are often contradictory goals since a better approximation of the multicast tree for improved performance demands for larger amount of information as compared to that available on the nodes, which in turn results in more overhead to adjust and maintain the tree with respect to changes in the network conditions and group membership dynamics.

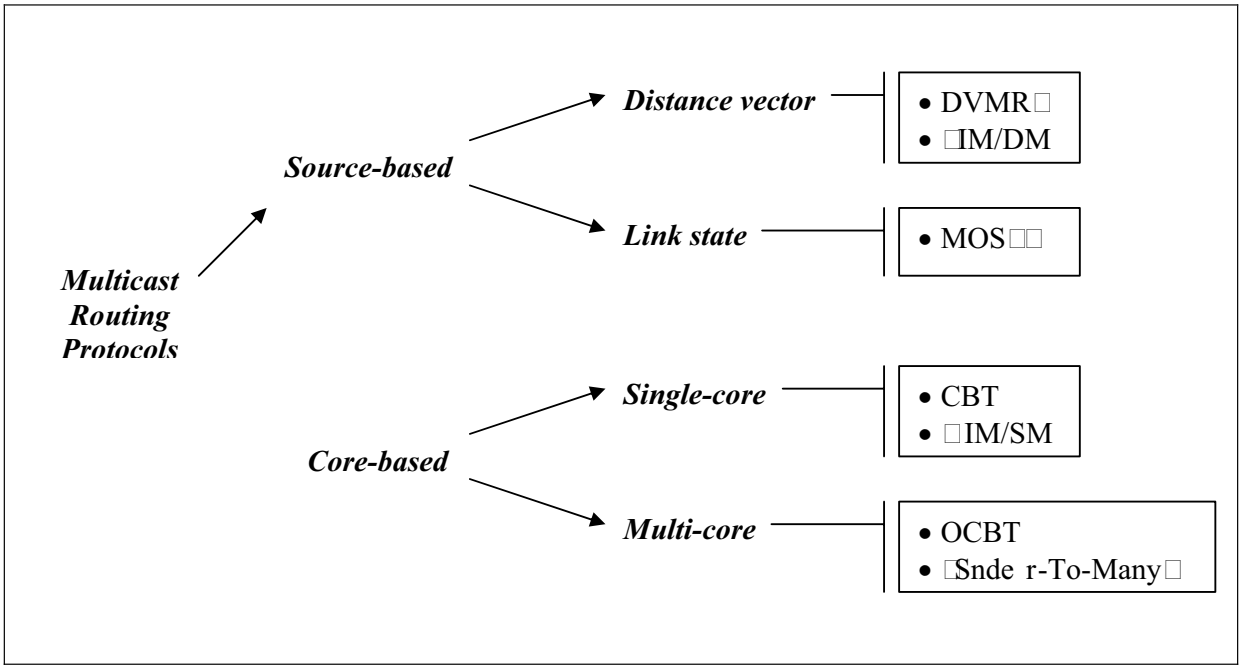

Figure 1. classification of multicast routing protocols.

In igure 1, we present a classification of multicast routing protocols with respect to their path construction and maintenance characteristics. The figure to a great extent also indicates the time-sequence of the emergence of the protocols with the exception of IM/DM which is the dense-mode component of IM protocol suite appearing in literature along with CBT. arlier protocols, namely DVMR and MOS assumed source-based multicast trees, i.e., the multicast path is built as a tree rooted as the source of the multicast group. The research later on shifted to core-based trees with the motivation of scalability in sparsely distributed multicast group in a wide domain. Recent trends in multicast protocols include multi-core multicast trees, which involve

more than one core in the multicast group for improved tree accessibility and performance. The assumed problem for tree construction in all multicast routing protocols in □igure 1 is S□T on delay metrics.

□.1 Source-Based □rotocols

In this group of protocols, the multicast tree is rooted at the source of the group with an attempt to minimi□e delay between the source and each receiver. □s we demonstrate in □igure 1, we further categori□e source-based protocols with respect to the unicast scheme they are built on.

Distance vector multicast routing protocol (DVMR□) [W□D88] is the earliest multicast routing protocol on Internet. It is an implementation of distance vector routing for multicasting and is an extension of RI□. DVMR□ is a rsr prooossus n.hussrrrvng

Source-based protocols in general are mostly suitable for small-scale, local-area applications. In sparse mode, the scalability of source-based protocols tend to degrade in terms of network resource consumption. □ major incentive in source-based architectures is the simplicity of the protocol as an extension of the underlying unicast platform. □s we will be reviewing in the following section, Core-based protocols consider the tree root not necessarily as the multicast source for scalability issues.

□.□ Single-Core Based □rotocols

The main incentive behind source-based S□T protocols is cost optimi□atio□ during multicast forwarding. S□T allows minimum available delay for forwarding the data stream to destinations of the receiver group. However, this comes at the expense of overall bandwidth consumption during transmission and the cost of protocol operation and maintenance.

Core-based trees separate the concept of source from that of the tree root. □ shortest path multicast tree is still constructed, but this time rooted at the core. However, the core(s) may or may not be the source. □ tree is formed still, rooted at a core this time. The choice of the core(s) can be made according to the topological location of the receiver group considering not only the delay in data forwarding, but also the bandwidth and protocol control overhead. In □igure 2, we present an extreme case when a core-based tree outperforms a source-based tree in terms of total bandwidth consumption, which is a significant advantage of the approach. The example includes a single source in the multicast group. The receivers are located close to one another in a topological neighborhood, and the source is far from their location. The dashed lines indicate a S□T rooted at the source minimi□ing the delay-distances to each receiver. The alternative is placing a core local to the receiver group and building the tree rooted at this core. The individual delays from the source to the receivers are not necessarily the shortest since each has to pass through the core. However, the data stream travels to the distant location of the receiver group only once, resulting in less bandwidth utili□ation. □or the same reason, core-based tree in this case is easier to maintain than a source-based tree.

There are two single-core based multicast routing protocols in the literature: □rotocol Independent Multicasting (□IM) [D□□+□6] and Core-Based Tree (CBT) [B□7, B□7b]. Both protocols build a single tree rooted at a node designated as the core. The core is advertised throughout the domain so that the new receivers to join the tree know whom to contact to initiate the process, and the sources to multicast the group know where to forward their stream for delivery. The multicast tree is constructed/updated as new members join the group.

CBT allows the bidirectional use of the tree so that when the data stream addressed to the core hits the tree at another node on its way, then that node acts as the tree root and relays the stream to every one of its links on the tree. Similarly, an on-tree CBT node, receiving data packets addressed to the multicast group it is serving, relays the stream to every tree link except the incoming one. The bidirectional nature of CBT makes more efficient use of the multicast tree in terms of the network resources. (Note that bi-directional trees are effective to serve this purpose in symmetric network environments. In asymmetric networks where the link cost and delay values are not necessarily equal in both directions, minimal cost is not guaranteed.)

□

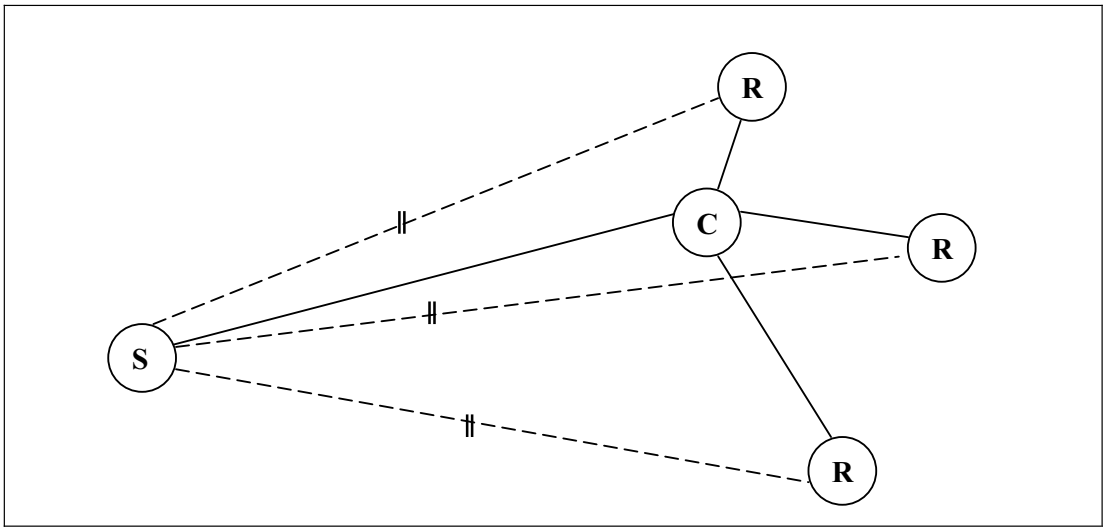

Figure □. □n example for core-based tree outperforming a source-based tree in case of a single source in the multicast group. The receiver group is located close to one another and far from the single source.

□IM is less reliant on the underlying unicast routing protocol in its operation. It refers only to the unicast state database for node distances. □ll message transmission, including control messages and transmission of the stream from the source to the core of the tree is governed by the protocol itself. □IM also offers individual receivers the option to receive the stream directly from the sender itself if the receiver decides that direct path from the source is more favorable for the path is shorter and/or the data rate warrants direct transmission.

Core-based trees are preferable to source-based trees in case of multiple sources in the multicast group. In this case, source-based schemes result in multiple shortest path trees with each having a source at the root, while a core-based scheme will have a single core based tree serving the entire members of the multicast group. The core-based approach then has the significant advantage of less control traffic of a single shared tree rather than multiple trees and, for the same reason, less state maintenance. Indeed a core-based tree can support a multicast session from more than one sender with only incremental increase in protocol overhead. Results in Wei and □strin [W□□□], and Calvert et al's [CZD□□] indicate that the advantages of core-based trees are eminent in sparse mode when the multicast group is distributed widely in a large portion of domain.

□.3 Multi-Core Based □rotocols

□ core on a multicast tree is a reference on the tree to the multicast tree for the non-member nodes. Since it is among on-tree nodes only the core(s) are advertised throughout the domain so that new receivers can join the tree and sources know where to forward the data stream for delivery. More cores would then mean more references to the tree and thus more efficient use of the tree. More reference points allow the selection of connection points to the tree according to

their proximity by the outsider routers, improving the resource consumption and traffic load. Multi-core tree schemes utili□e this idea under distance-vector architecture.

There are different mu

less likely that the entire traffic will be concentrated around one particular core in the multi-core approach. In single-core trees, traffic concentration is high around the core especially in case of multiple senders in the group. In multi-core architectures, the traffic concentration decreases as the number of cores increase especially when the cores are randomly distributed, since the traffic is spread around the domain on different cores. [CZD□ □, Z□L□2]

Zappala, □abbri and Lo's results [Z□L□2] indicate that multi-core trees with effective sub-optimal core placement mechanisms result in less delay than single-core based trees with optimum placement of the unique core, since the multi-core case allows trees which are likely to be distributed adequately to be closer to a group of sender(s) than a single core. Similarly, multi-core trees incur less total cost of delivering the multicast stream in terms of bandwidth utili□ation than a single core tree since a new node to join the tree is likely to locate a core among the alternatives, which is potentially closer to it than the only core alternative which is the tree root in the single-core based tree.

3. Core Selection Algorithms

Core selection is the problem of selecting the placement of a core or cores in the domain for the purpose of improving the performance of the tree(s) constructed around these core(s) and thereby the performance of the multicast routing protocol according to the predetermined metrics. In this section, we analy□e core selection algorithms in the literature and classify them with respect to certain characteristics based on their deployment. In the following subsection, we provide notations and definitions In Section 3.2, we present our classification scheme. We review the core selection algorithms in Section 3.3. □ollowing this review, in Section 3.□, we analy□e computational and message complexity of each algorithm. Section 3.□ includes further discussion on core selection methods.

3.1 □reliminaries

Before we proceed, we review in this section the terminology and preliminary research related to core selection. Wall's study [W8□] on single core selection problem for broadcasting is precedent in this line of research in the literature. The terminology we introduce is parallel to his, and adequately covers the generic terminology we use throughout our review since Wall's study in context of broadcasting outlined this line of research, and is the foundation of early contributions to core selection for multicasting.

We denote minimum delay distance of any given node pairs *u*, *v* in the domain by *d*(*u*,*v*). □or a shortest path tree T_c rooted at node *c*, D(s,T_c s ndhgroupbrsonghr T_c D(s,T_c D(c,T_c d(s,c) D(c,T_c T_c □(s,Ts□(

minimum among all such possible trees. *diameter-centered* tree is the tree with the smallest possible diameter. , , ,
ofvrgdysvrgdyndvrg ofvrdys icaetaie
sss sss

multicasting via a core-based tree of which the core is selected to yield a diameter-centered or *D*-centered tree.

However, Wall's [W8□] findings are for broadcast environments and their applicability for multicasting are limited with the necessity of considering the sources and all candidate cores as on-tree alongside the host group members in the calculations for core selection. With this in mind, Wall's findings suggest that although *A*-centered trees may have higher upper-bounds on certain measures, they are more consistent to the average performance than the other types of trees.

3. □Classification of Core Selection Algorithms

Before proceeding with the review of core-selection schemes, we present a classification for their deployment and structural and operational characteristics (□igure 3).We mainly distinguish core selection algorithms as *single-core* versus *multi-core* for this is the dominant determinant of the complexity of the scheme. □arallel to the emergence of single-core based shared trees before multi-core based architectures, research on single-core selection precede multi-core selection in literature.

Algorithm	***M***	***C***	***A***	***D***	***Metrics***
maxDist [TR□7]			+	+	hop-count
aveDist [TR□7]			+	+	hop-count
Diam [TR□7]			+	+	hop-count
estCost [TR□7]			+	+	hop-count
Topology-based [CZD□□]			+	+	hop-count

Group-basdnr

The underlying platform assumed in all core selection algorithms is the distance-vector routing environment in which restricted domain information, namely the distances from the router itself to each other domain router are available to a particular router. In our context, we assume the availability of this information in every metric. □ccording to this, the state information of a given router includes the shortest delay, cost and hop-count distances from itself separately to every other domain router along with the outgoing links to lead to the shortest-distance path.

The extent of the use of local information on domain nodes for the operation of the algorithm determines whether the algorithm is centrali□ed or distributed. □ *distributed* algorithm is one that operates on multiple nodes using the local information on each node. The solution is obtained as the nodes are traversed, or as the performance results obtained locally from each candidate are gathered at one node by message transmission. In a *centralized* algorithm, the entire processing is carried at one node, and the information required for the process is forwarded to it. We classify an algorithm as centrali□ed if its deployment to operate on one router is more feasible than its distributed implementation, distributed otherwise. Once again, the assumed state information available to all routers in the domain is the distance-vector environment□ the distances in each metric from the node that the state table resides on to every other router. □s we analy□e the algorithms for their computational structure in Section 3.□, we further distinguish a group of distributed algorithms for their operational characteristics as □Compute-and-Report□ *Class-CR* algorithms.

We denote a core-selection algorithm as *constrained* if the algorithm considers a given delay-bound for the group during the selection process, and the resulting core set is such that there exists a path between each source-receiver pair in the group which passes through at least one core without violating the delay-bound.

□n *asymmetric* network is one in which the link costs in all metrics are not necessarily equal in both directions. It is clear that the algorithmic complexity reduces when the network is symmetric especially in distributed processing. This is because in the assumed distance-vector setting, the information available to each node represents the distance information not only
ssss

nwork, hv wdr ppby.

gorhovrpwh
p onsnonprossoprng.hons

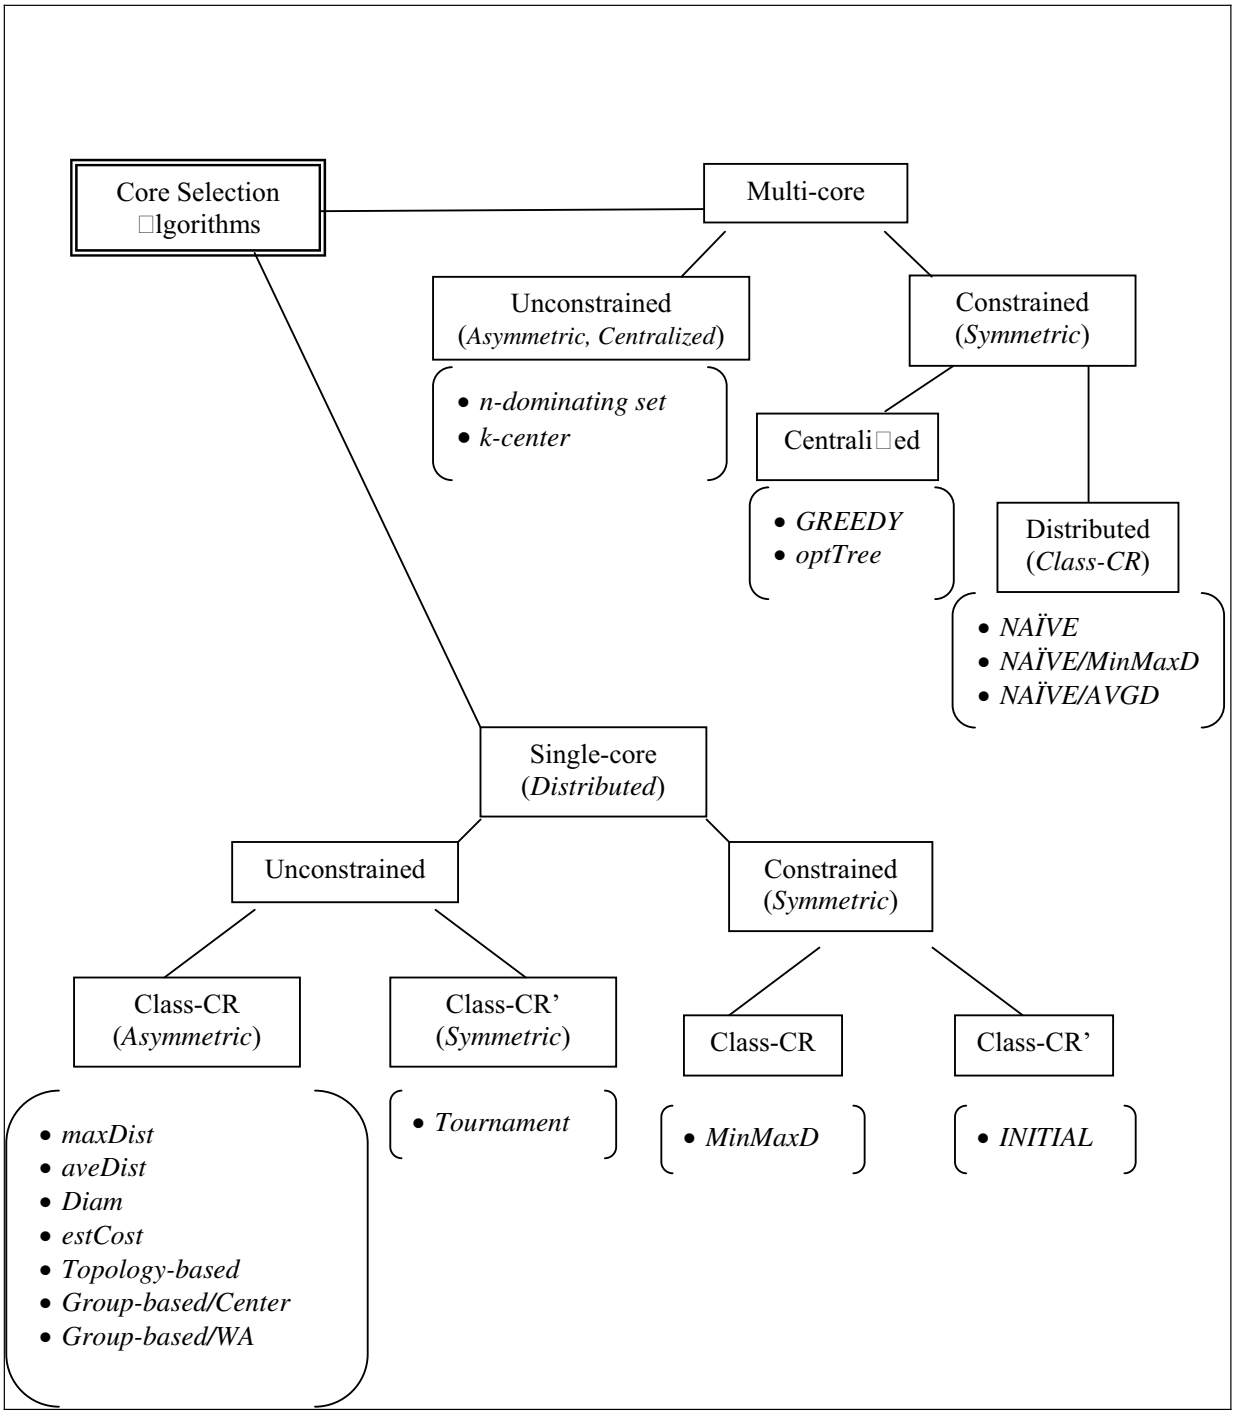

Figure 3. Classification of core selection algorithms. Class-CR' denotes those distributed algorithms which are not in Class-CR.

3.3 Review of Core Selection Algorithms

In this section, we review the core selection schemes according to the selection criterion they apply and their performance. Our analysis in Section 3.□ completes the characteri□ation of these algorithms for their computational properties.

3.3.1 Single-Core Selection Algorithms

□s previously stated, earlier protocols applied a single-core scheme for the ease of construction and maintenance of a loop-free multicast path. The relative simplicity of the selection of a single core to serve the entire multicast group allows feasible deployment of these algorithms on local information and moderate message exchange among the participant nodes.

3.3.1.1 Single-Core Selection Algorithms / Unconstrained

Thaler and Ravishankar [TR□7] adapt Wall's [W8□]broadcast single core selection criteria to multicasting to a subset of nodes in the domain. □n additional selection criterion, estimated cost, is also introduced. The estimated cost is the average of minimum and maximum possible tree costs computed. The applicability of the results is restricted to the case where all the link costs in the network are assumed to be equal. However, the core selection algorithms in [TR□7] apply for asymmetric networks for they rely on local information on the nodes during the selection process. The criteria considered are:

1.) *maxDist*: maximum distance from the core to any multicast group member. Corresponds to *D*-centered trees described by Wall [W8□]for broadcasting environment.
2.) *aveDist*: average distance from the core to all multicast group members. Corresponds to Wall's *A*-centered trees.
3.) *Diam*: diameter of the tree rooted at the core defined as the sum of the distances of the two farthest away members in the group from the core. Note that this definition does not describe the exact diameter of the tree rooted at the core minimi□ing this measure since the diameter of a shortest path tree does not necessarily pass through the root [W8□].
□.) *estCost*: the average of the minimum possible and maximum possible cost of the tree rooted at the candidate node being tested. The minimum cost is estimated to be the sum of the maximum destination cost and the number of duplicate distance nodes in the multicast group, on the assumption that the minimum cost tree is of a chain structure, with the exception of the group members that are of equal distance to the core with one another, adding minimum of one more link each to the tree. The other extreme, tree of highest cost, is of star structure with all multicast group members directly connected to the core. The exception arises when the degree of core is less than the number of multicast group members, in which case multiple group members share the links outgoing from core. The minimum estimated cost, $estCost_{min}$, and the maximum estimated cost, $estCost_{max}$, are given by:

$$estCost_{min} = \max_{u \bullet S} d(\quad s \; s \quad)$$

$$estCost_{max} = \begin{cases} \sum_{u \in S} d(c,u) & \text{if } |S| \quad degree(c) \\ \sum_{u \in S} d(c,u) - (|S| - \deg ree(c)) & \text{otherwise} \end{cases}$$

$$estCost = (estCost_{min} + etCt \qquad min$$

c shnddor,nd *S s h us group bng onsdrd for vuon*

od

uretesting

distance to any other node in the graph is minimum possible across all nodes in the graph. Consider the function (u,Tv),dfonofhfunon D(u,Tvp ratherthandelay-ditancebeteenthedom aieicae oadcating.Then,anode u (u,Tu min vVV (v,Tv V D a

single and multiple sources in the group, with the difference in this cost gain diminishing as the number of sources increase and the network is highly connected.

3.3.1. Single-Core Selection Algorithms / Constrained

Salama [Sa96] studied two single-core selection algorithms for delay-constrained shared trees in symmetric networks. Salama's solutions are restricted to the case in which all receivers are also sources. A typical example of this case is a video-conference application, where all participants can communicate with every other member in the group both ways. Salama tested the following delay-constrained single-core selection methods:

1.) *MinMaxD*: the node that has minimum of maximum delay distance to any multicast group member. This corresponds to the delay-constrained extension of *MaxDist* proposed in [TR97].
2.) *INITIAL*: A random multicast group member is selected to initiate the core selection process. The selected member locates the most distant group member from itself, and constructs a delay-constrained path in between. The node closest to the center of this path is selected to be the core. *INITIAL* applies the main intuition behind the *Tournament* selection algorithm [SBK92] to a lesser extend, considering the delay constraint between only two distant members of the multicast group. *INITIAL*, however, does not guarantee success to locate a core that would result in a tree spanning all multicast group members within the delay bounds, even if such a core existed.

The delay bound to be fed as input to the delay-constrained tree construction algorithm varies depending on which method is used for the selection of the core. Note that with Salama's assumption that all multicast group members are both sources and receivers, the delay-constrained tree construction problem transforms to the construction of a "diameter-constrained" tree, in which the longest on-tree path passing through the core is within the given bound. With this, Salama suggests that the delay bound, Δ', used for *MinMaxD* is at most as great as $\Delta/2$ where Δ is the actual delay bound for the group to ensure that an on-tree path crossing the core between a source-destination pair does not exceed the actual constraint. For *INITIAL*, the on-tree delay constraint value to be considered, however, differs in that the delay from the core to one of the two members on the initially constructed path can exceed half the value of actual value. The on-tree delay constraint value for this latter case, therefore, is computed as $\Delta' = min(\Delta/2, \Delta - max(d_1, d_2))$ where d_1 and d_2 are the delays on paths between the core and the initial multicast group members joining the tree to select the core.

Both selection criteria are based on a single metric – delay. Cost of the trees in both cases improved closer to the cost of the optimal tree as the number of group members increased, verifying the fact that the cost-performance of the tree is higher in dense mode. *INITIAL* outperformed *MinMaxD* in terms of cost especially when the delay bound is tighter. Note that, this result is biased by the fact that the tree cost increase is based on successfully generated trees satisfying the delay bounds by the use of a polynomial-time tree construction algorithm, which lacks well-approximation of cost minimization. Performance results indicate that the difference

in costs of the trees generated in *MinMaxD* and *INITIAL* remained within 1□% of the cost of the optimal tree, with this difference diminishing as the number of group members increase.

3.3.□Multi-Core Selection Algorithms

Multi-core algorithms have the common property that their feasibility is eminent in centrali□ed implementations using extensive domain knowledge compared to the local state information on distance-vector protocol routers.

3.3.□.1 Multi-Core Selection Algorithms / Unconstrained

Zappala, □abbri and Lo [Z□L□2] tested two multi-core selection algorithms against average core choice in *Senders-to-Many* and *Members-to-Many* protocol settings:

1.) *n-dominating set*: locates a minimal dominating set as the set of cores, which is a subset of nodes in the domain so that all nodes in the multicast group are within *n*-hops of this set for a given *n*.
2.) *k-center*: places a given number of *k* cores in the domain so that the hop-distance from all nodes in the multicast group to the cores is minimum.

Note that neither the dominating set nor *k*-center algorithms has polynomial-time solutions [G□□□]. Zappala et al [Z□L□2] also use the greedy dominating set algorithm with modifications for *k*-center heuristics. □ node *r* is said to be dominated by a core *c* if *n* is within at most *n*-hops distance of *c*. We denote by *domination set* of a given node *v* the set of nodes that are dominated by *v*. Dominating set, on the other hand, refers to a set of nodes so that the set of receivers is a subset of the union of domination sets of the members of the dominating set. The algorithm iterates to select the node in domain that currently dominates the maximum number of un-dominated members until all the group members are dominated. The dominating set heuristics is further parameteri□ed with a constant *k* to specify the exact number of cores to be selected for a solution to *k*-center problem. It should be noted that Zappala et al [Z□L□2] were not specific on the algorithmic details of *k*-center heuristics.

Zappala, □abbri and Lo [Z□L□2] tested each of these algorithms comparing their performance to random core selection. Both *k*-center and dominating set algorithms demonstrated modest improvement in delay compared to random core selection. *k*-center core placement with multiple cores resulted in lower delays than single-core tree with optimal core placement as the value of *k* increased. Zappala et al's results suggest that polynomial-time heuristics for dominating set, and especially *k*-center core placement contribute to the efficiency of multi-core protocols.

Both *k*-center and dominating set algorithms refer to hop-count as the single metric for core selection. We assume, due to the complexity of especially the *k*-center algorithm for potential heuristics, that both algorithms are implemented as centrali□ed, since the transmission of relevant domain information to the operating node in centrali□ed deployment is likely to match for the advantages of the distributed implementation which involves excessive message exchange among the participant nodes for the same level of approximation of the optimum result as a

centralized algorithm provides. Hop-count is an asymmetric metric for it describes equal link weights in both directions by definition, and hence *k*-center and *n*-dominating set algorithms operating on hop-count classify as asymmetric.

3.3.□.□ Multi-Core Selection Algorithms / Constrained

Salama [Sa□6] also studied the problem of multiple core selection for Zappala, □abbri and Lo's [Z□L□2] Senders-to-Many architecture for delay-constrained environment. Observe that, for delay-constrained case, the additional motivation for multi-core as opposed to single-core scheme is that multiple cores extend the potential solutions to delay-bound paths between source-receiver pairs passing through a core where a single core could fail to serve the entire group in this manner without violating the delay constraint. Salama's multi-core selection algorithms assume a symmetric network:

1.) *GREEDY*: This algorithm uses a variation of the solution for the dominating set problem in which the domination criterion is the delay bound rather than the hop-count. Note that the problem remains N□-complete in the modified version [Sa□6]. *GREEDY* assumes as input, for each node in the domain, the set of receiver nodes that the node is domi□ating for the delay bound for all sources in the group. The algorithm merely iterates to select the node that dominates the maximum number of un-dominated receivers as the next core until all receivers are dominated. *GREEDY* operates on full domain information and is centrali□ed.

2.) *NAÏVE*: *NAÏVE* applies iterative selection of the core set as the tree is constructed. Initially the selected core set is empty. When a new receiver attempts to join the multicast path, the algorithm checks iteratively the members of the core set to see whether there already exists a core so that the paths from each source to the new receiver via the core is within the delay bound. If the search fails, then the candidate core set is searched for a core to lead to a feasible solution. Salama suggests as the candidate core set either the set of receivers, or an ordered list of a set of nodes with respect to the maximum (*NAÏVE-MinMaxD*), or average (*NAÏVE-AVGD*) delay of the candidate to the multicast group members. When the receiver set constitutes the candidate cores, the next core to be selected is always the new receiver to join itself i□ case of inexistence of a feasible core among the selected ones. In all versions, *NAÏVE* processes for at most as many cores as the number of sources in the group be selected, and terminates without providing a solution with the idea that source-based trees would be more feasible when the number of cores and hence number of distinct trees exceed the number of source-based trees. *NAÏVE* and its variations operate on the minimum delay distances between source-candidate core and receiver-candidate core pairs which is maintained in the routing state database of these nodes. Thus, *NAÏVE* can be implemented for distributed processing alternatively with the transmission of the delay-routing state information of the group nodes and candidate cores to a coordinator node in the domain to execute the algorithm, or with the transmission of group information to the candidate cores where each candidate evaluates itself against the group members and returns its result back to the coordinator for comparison with the other candidates. We will be comparing commenting these two implementations in Section 3.□ during our complexity analysis.

Both *GREEDY* and *NAÏVE* return a feasible solution whenever such a solution exists, provided that the candidate core set in *NAÏVE* does not exclude some subset that can actually lead to a

feasible solution. Both algorithms lack the consideration of the cost within the core selection criterion. In terms of the number of cores generated, *GREEDY*, considering the delay-domination of the group members directly throughout the algorithm, performed very close to the optimal solution. *NAÏVE-MinMaxD* and *NAÏVE-AVGD* underperformed *GREEDY*, generating similar results still close to optimal.

Karaman and Hassanein [KH□3] extend *GREEDY* to consider an □optimi□ation criterion□ in addition to the □domination criterion□, which still is the given delay-bound between any source-destination pair in the group. The algorithm, which we call *optTree*, generates a set of cores each associated with a domination set □ the receiver partition including those receivers dominated by that core for all sources in the group. The assumed design problem for *optTree* is the delay-constrained minimi□ation of total number of links on the resulting set of paths for the multicast group for optimi□ation of tree structure on total link count. During its iteration to extend the selected core set until all receivers are dominated, *optTree* updates the existing domination sets according to the optimi□ation criterion to better approximate the optimum set of multicast paths to be constructed. □ccording to this, each receiver dominated by the new core and in the domination set of the already selected core is assigned to the domination set of that core ranking higher with respect to the optimi□ation criterion. The following are considered for optimi□ation criteria:

1.) *Node-count*: highest number of nodes that the node is dominating
2.) *Delay-residue*: highest average delay-residue from the core to the receivers in its domination set which are not yet dominated by a selected core
3.) aveHop

□.) _____

b

opTrNod-ount

triggering of participant nodes to mark the beginning of the process, and ends with the set of core(s) selected among the candidates. We exclude from this cycle the announcement of the selected core(s) to the domain. □s we proceed in this section, we will be observing that for some of the algorithms, the candidate core set□the pool of nodes among which the core-choice is to be made□ can be restricted externally to the selection algorithm without affecting the modular operation of the algorithm. In such cases, we assume as input to the core selection process the set of candidate cores, and exclude the process of restricting this set from core selection. □or distributed algorithms, we generically name the node designated for the gathering of local results from the participant nodes, and processing on this compiled information, where necessary as the *coordinator* node. In the case of *Tournament* it is the □scheduler□[SBK□□]. In the centrali□ed case, we use the term □coordinator□to represent the only processing node in the domain. We assume that the coordinator node is already selected and known in the domain.

Algorithm	***Local***	***Local-nodes***	***Coordinator***	***Message***
maxDist [TR□7] (*)	M	COCOC		

the entire algorithm. Note that, in the centralized case, only the "Coordinator" column of the table applies and the rest of the row entries are nil. The assumed state information available on a given router in all cases consists of the distance information from the router itself to every other router in domain in each relevant metric. For centralized algorithms, we further comment in Section 3.1 on the amount of message transmission to communicate the required information to the coordinator under the availability of restricted domain knowledge as assumed.

We note that a number of a distributed core selection algorithms have similar computational characteristics to locally compute to evaluate the candidate cores and report the results to coordinator for comparison. We group these in a class and denote it as the class of "*Compute-and-Report*", *Class-CR*. In Class-CR algorithms, each candidate node measures itself for the selection criterion based on the local information and reports the result to the coordinator. The coordinator compiles the node results from the candidate cores and compares them to select the best scoring node as the core. All selection algorithms proposed by Thaler and Ravishankar [TR97] and Calvert et al [CZD+], and *MinMaxD* single core selection algorithm [Sa96] are in this class. Message exchange involves the transmission of the group members information from the coordinator to the candidate cores, and the node-results back to the coordinator with the exception of *Topology-based* core choice [CZD+] in which each node testing itself against the values for all domain nodes, the communication of group information to the candidate cores is eliminated. Therefore $O(|C|)$ messages are transmitted for the selection of a particular core in Class-CR algorithms, where C is the candidate core set. Selection of the best core among the candidates on the coordinator node is done by iterating over individual node values with constant-time computation at each iteration, and requires linear time in the size of candidate core set provided that the node results are transmitted to coordinator within a given time-out period. The time-complexity of computation on coordinator is the same as the bound on number of message exchange for all Class-CR schemes.

Local computations of Class-CR algorithms iterate for each node in the multicast group to calculate the node measure on the selection criteria which, in each case, requires to look up the local state table followed by the calculation and comparison with current optimum in constant time. Thus, local computations in Class-CR algorithms take $O(|M|)$ time where M is the set of multicast group members. Specifically, observe that *MinMaxD*, selecting the among the candidates according to the same criterion as *MaxDist* unless no node among the candidates satisfying the delay constraint exists is computationally same as *MinMax* with the resulting core additionally tested for in constant time to satisfy the delay bound. Similarly, *Topology-based* core choice operates the same algorithm as *MaxDist* with the difference that the "multicast group" this time is the entire set of domain nodes, and the link weights are unit for hop-count metric. For *Group-based/center* choice of a core, a similar process takes place, but this time on the modified graph of which the nodes are the receivers and weight of each link connecting a pair of receivers is the total hop-count of the minimum hop-distance path connecting the same two receivers in the original graph. Note that *Topology-based* core choice, operating solely on the domain topology stands out among all others in that it requires only static processing irrespective of the varying multicast groups active in the domain.

If we consider *NAÏVE* operating in a relatively static network environment where routing state updates are unlikely during the multicast session, one-time transmission of multicast group and

candidate core delay-routing state tables to the coordinator and maintenance of this information at the coordinator allows local operation for the selection of core set. However, considering the dynamisms of network environment, we assume the core selection algorithm executing on up-to-date network state information for the selection of each core. □ccording to this, we assume the distributed design of the algorithm so that, when a new group member request to join, the coordinator transmits the identification of this node to the candidate cores. □ach candidate operates on its routing state database to test itself for the evaluation criterion depending on the *NAÏVE* version, and returns its result back to the coordinator. Thus, at each iteration of *NAÏVE* for the selection of another core, the algorithm operates as a Class-CR algorithm. The message transmission involves, as in other Class-CR algorithms, the transmission of group information to the candidates and the node results back to the coordinator and hence is bounded by $O(|C|)$. □or the selection of a single core, the algorithm operates in $O(|C|)$ time at the coordinator for message transmission and comparison of node results. During a local execution, each candidate core processes on its distances to the group members and thus the complexity of computation during a core selection is $O(|M|)$. *NAÏVE* iterates this cycle at most as many times as the number of group members. Therefore, the message complexity, and the complexity of local computations at the coordinator and the candidate nodes are multiplied by $|M|$. The local processes execute at most as many times as $|C||M|$ where the operating local-node set is M in *NAÏVE*, and C in *NAÏVE*/variations.

Tournament suffers from heavy message exchange. □ach level of tournament is the sequence of following steps:

b

d

hrndoprngprossnsp(ib

tree, and the number of pairs depicted on it is $|M|$-1. When one member is removed from the set M, one competing pair involving the removed node turns into an unmatched node at the first tournament level due to an odd number of nodes at that level. Similarly, the removal of each further node from M results in the loss of exactly one pair in one level further set of pairs until exactly one half of the binary tree is lost and the multicast group size is $|M|=2^{n-1}$. The number of times the midpoint-computation process is executed is exactly the number of pairs matched throughout the entire tournament, and is exactly one less than the group size with this argument.) As we stated above, each execution of the local process at step (*iii*) requires message transmission of $O(|V|)$, hence the overall number of messages at this step is bound by $O(|V||M|)$. Additional message transmission which take place at steps (*ii*) and (*iv*) in Tournament amount to $O(|M|)$, and total number of messages exchanged throughout the $O(|V||M|)$.

INITIAL operates similar to *Tournament* at a single step in that one randomly selected group member (*i*) determines the member farthest from itself in delay, and (*ii*) locates the domain node closest to the "midpoint" of the shortest-delay path between the two. Using a similar argument as in *Tournament*, steps (*i*) and (*ii*) are executed respectively in $O(|M|)$ and $O(|V|)$, resulting in time complexity of $O(|V|)$ of the algorithm. The number of messages transmitted in step (*ii*) is bound by $O(|V|)$ and constant in step (*i*).

NAÏVE iteratively adds a new core to the set of selected cores when none of the current cores satisfy to serve a new multicast group member within the given delay bound. For each new group member, the existing core set is searched for a satisfactory core. During this search, each core is examined whether to meet the delay bound for each source in the group to transmit the receiver. Assuming knowledge of the farthest source in delay-distance to each selected core is maintained throughout the algorithm, the time complexity for searching the current core set is $O(|C_{select}|)$ where C_{select} is the set of selected cores and is bound by M the set of multicast group members. If the search in C_{select} fails, then the new receiver itself is designated as a new core provided that it is in reach within the delay constraint from every source in the group executing in time linear to the number of sources in the group. Therefore, the entire algorithm computes in time $O(|M|^2)$. In case of *NAÏVE/MinMaxD* and *NAÏVE/AVGD*, the candidate core set, which is the set of domain nodes, is ranked respectively for the maximum and average node delays to the group members. This process executes modularly to the generic NAÏVE algorithm with the modification of *NAÏVE* to operate on an input candidate set, and takes $O(|M||C|)$ time. The computational complexity of variations of *NAÏVE*, therefore, is $O(|M|^2+|M||C|) = O(|M||C|)$.

A greedy dominating set heuristic iterates for the following steps until all the multicast group members are dominated by the core set:

i.) select the candidate core, *c*, ranking highest according to the "optimization criterion"
ii.) add *c* to the set of selected cores
iii.) update the domination sets.

As mentioned earlier, we refer by *domination set* to a set associated to a candidate core to return the set of currently un-dominated multicast group members. A typical implementation maintains a two dimensional table with the multicast group members and the candidate cores being separately the indices of these dimensions. Note that, in case of delay-constrained

implementation of heuristic, the process to generate the domination sets considers the delay-distances separately between source-core and receiver-core pairs in two different cycles without extending the computational complexity of the process. The greedy algorithm needs as input the set of domination sets, which can be constructed in $O(|M||C|)$ time for all candidates. Updating the domination sets in step (*iii*), following the same procedure on already-constructed sets, have the same time complexity. Step (*i*) operates on domination sets to count the number of members in each set to find the one that has the maximum set size, executing in time proportional to $|M||C|$ in the worst case. Step (*ii*) clearly executes in constant time. Steps (*i*)-(*iii*) are iterated exactly once for each group member in the ultimate case. Therefore, the complexity of the entire algorithm is $O(|M|^2|C|)$. The generic algorithm we described applies for both *n*-dominating set problem [ZL02] and *GREEDY* [Sa96] with modification on the initial construction of domination sets so that a group member is in domination set of a candidate core if it is in respectively within the *n*-hops or delay-constraint's reach from that candidate.

The optimization criteria suggested for *optTree* [KH93] can be computed for one core in the iteration that the core is examined as a candidate, without changing the structure of the greedy algorithm. If we consider a combination of *Node-count*, *Delay-residue, aveHop* and *Degree* as the optimization criterion of the algorithm, the additional computational complexity to the greedy algorithm is clearly that of the one among four with the highest complexity. In *Node-count*, *Delay-residue* and *aveHop* of *optTree*, each candidate core is examined for its rank depending on the receiver set they are potentially dominating, repeating at most as many times as the number of receivers. *Delay-residue* and *aveHop* also examine the distances of the core from the sources for residue in delay and hop-distances to the receivers. Therefore, the computational complexity of *optTree* is $O(|M|^3|C|)$.

As we mentioned in Section 3.1, Zappala et al [ZL02] suggest a modification of the dominating set approximation algorithm for a *k*-center heuristic. Assuming repeated execution of *n*-dominating set heuristic for increasing values of *n* until a set of cores of maximum size *k* is obtained to dominate the all members of the multicast group, the complexity of *k*-center approximation for algorithm core selection is $O(|M|^2|C|h_{diam})$ where h_{diam} is the maximum hop-distance considered between a core and a group member and is bound by the diameter of the domain in this metric.

3.4.1 Implementation of Centralized Algorithms in Distributed Environments

As mentioned earlier, local information available to routers in generic setup is the distances of the router the state table is residing on to every other router in the domain in all metrics. Centralized algorithms are not readily supported in this setup for these algorithms demand and operate on broader domain information then that locally available. Our classification of centralized algorithms is based on their higher feasibility compared to their distributed implementation on multiple nodes with message exchange among the nodes, as opposed to processing on one particular node with the required domain information transmitted to it in centralized version. We further examine the amount of message transmission in centralized algorithms to facilitate their deployment in distance-vector routing environment.

We examined the generic dominating set heuristics as iterative steps, for each new multicast group member, of selecting the best candidate core to serve the next new member of the multicast group and updating the domination sets accordingly for the next iteration. If we look into these steps, they all operate on the distance-vector information of the candidate cores to determine the domination set of each candidate for a given multicast group, and examine each candidate with respect to □optimi□ation criterion□. Then, an algorithm to approximate the domination set solution following this heuristics simply transmits the distance-vector information of each candidate core in all relevant metrics to the coordinator. Similarly, the coordinator is notified of new members to join the multicast group merely for their domain identification. Therefore, assuming a particular candidate core set and stable network conditions, each candidate transmits its state table once to coordinator in $|C|$ messages, and the coordinator is notified of the members of the multicast group in $|M|$ number of messages, resulting in overall message complexity of $O(|M|+|C|) = O(|C|)$ on assumption that $|C| \geq |M|$.

3.□.□ Bidirectional Im□lementation of *GREEDY* and *optTree*

In a *bidirectional* routing tree, the first on-tree node receiving the data stream acts as a tree root and the stream is forwarded from this □root□ towards all leaves. If the tree is not maintained bidirectional, the stream travels first to the root of tree and then delivered throughout the tree.

Whether the multicast tree is maintained as a bidirectional in a symmetric network environment is especially significant in constrained dominating set heuristics, since a bidirectional tree can describe shorter-delay paths between source-receiver pairs than a non-bidirectional one on which every path between such a pair has to pass through the tree root. □ach core has broader potential domination in this case, resulting in better approximation of optimal core selection. The core selection process in this case, however, demands not only for the evaluation of delay-distances between multicast group members and candidate core sets, but also the corresponding information between the members of these sets and any other domain node, for any such node is a potential intermediary on the tree to be developed to act as the root for certain data source(s) on the tree. So, the bidirectional consideration of the multicast tree to be governed by the core set is a trade-off between the optimality of the result and the complexity of the core selection process.

Salama [Sa□6] approaches this problem by first computing the domination set of each potential core before proceeding with the iterative selection process to produce a set of nodes to dominate the group. During the computation of the domination sets, each source, receiver and potential core, (s,r,c) triple, is examined for a delay-bound path between s and r on the core-tree rooted at c. Since the tree to be constructed is assumed to be bidirectional, the path between s and r passing through c can be shorter than the concatenation of the paths between (c,s) and (c,r) pairs if these two paths overlap partially at the core end. In other words, there may exist a node v which is on the shortest paths between both the pairs (c,s) and (c,r), and thus the portion between c and v is shared by both paths. The minimum-delay between s and r via the tree governed by c in this case reduces from the sum of the delays between (c,s) and (c,r) to that of (v,s) and (v,r). The corresponding computation to generate domination sets examines each node separately between (c,s) and (c,r) node pairs to locate a overlapping paths. The process then requires as input the knowledge of the entire domain nodes, and the message transmission for the bidirectional implementation of constrained

dominating set algorithms is $O(|V|)$. As stated by Salama [Sa96] the process iterates for each source, receiver and potential core triple on all intermediary nodes lying on the examined paths, and the computational complexity of the preliminary process for the generation of domination sets is $O(|M|^2|C||V|)$. The additional component of both algorithms to select the cores among the candidate pool under this approach is computationally less complex than the preliminary phase. Therefore, the complexity of *GREEDY* [Sa96] and *optTree* [KH93] on assumption of bidirectional trees is $O(|M|^2|C||V|)$. However, the effectiveness of bi-directional design of core selection algorithms is restricted to multicast groups in which the source and receiver sets overlap to share paths on the core-trees.

3.5 Discussion

The integration of core-selection algorithms into core-based multicast routing protocols is dominated by whether the components are serving single or multi-core scheme. the single-core selection algorithms are restrictive for the performance of multi-core protocols and the algorithms to generate multiple core can not co-operate with single-core protocols. The remaining characteristics of core-selection algorithms superior to the protocols clearly do not restrict their applicability in coordination with these protocols.

Hop-count assumed as the only metric in most core selection algorithms restricts the potential results of the implementing protocol for all protocols described use delay-distance for exact measure. Hop-count, however, can be replaced by delay as the metric in the algorithms proposed by Thaler et al [TR97] except for *estCost*. in *maxDist*, *aveDist* and *Diam*, hop-count is used as a measure for point-to-point distance and can be replaced by delay without affecting the algorithm structure. In case of *estCost*, hop-count is measuring the overall tree structure and is the metric in context. Likewise, the description and measurement of topology center can be extended by allowing variable link weights during the process for computing the topology and group center. modifying *Topology-based* and *Group-based/Center* [CZD94] core selection algorithms for delay metric. The problem of *k*-center is defined for optimization on average hop-distance of the multicast group to the core set for a given size of core set, and can be modified as a problem to optimize on average delay-distance with an extended heuristics. *n*-dominating set problem has a different characteristics in that the hop-count is used as the domination criterion to be referred to rule in/out domain nodes into the candidate-core pool to be further filtered by the optimization criterion which in case of the applied heuristics by Zappala et al [ZL02] is the number of multicast group members dominated by the particular node. A certain delay-distance to be given as the dominating criterion rather than the hop-distance imposes delay as the more precise metric during the selection process. However, the delay-constrained dominating set algorithm, namely *GREEDY* [Sa96] readily provides a broader solution in terms of potential performance by considering the given delay bound from source(s) to receivers via core(s) rather than separately between source-core and receiver-core pairs with a process of similar complexity. Cost in *Tournament* [SBK94] can be replaced by delay with no change in its algorithmic structure since both metrics are additive and have similar characteristics. Note that hop-count mainly serves for measuring the tree structure through measuring the number of on-tree links as it is perceived in [KH93] as additional metric to delay, and clearly has positive correlation with delay.

The core-based protocols studied so far assumed symmetric networks only for they all initiate the path construction process at the new member to join and measure the path distance from receiver to core rather than the other direction in which the stream flows. The asymmetry in path construction and maintenance can be accomplished by delegating the path construction to the core from the receiver-to-join simply by traversing the path to be added to the tree as the acknowledgement travels towards the receiver□ rather than as the join request travels the core. □roviding asymmetry in protocol features to reali□e the tree structure is far less challenging than the algorithmic components□ including core selection□ to specify the tree structure, and existing protocols apply with straightforward modifications to asymmetric networks. Note here that, hop-count is an asymmetric metric and characteri□es the algorithm asasymmetric regardless of its symmetry characteristics. Those distributed algorithms we concluded to be applicable also for delay metric above, however, maintain their asymmetry since they all utili□e the link values in their exact direction during the computations.

□ll core selection methods proposed by Salama [Sa□6] are constrained, designed to meet the given delay bound on the tree to be constructed rooted at the selected core(s). The core selection algorithms studied by Salama [Sa□6] can be described into protocol suites by their extension with tree maintenance and failure recovery components. Salama is leaving the cost optimi□ation to the phase of tree construction and is not considering it as an additional design goal for core selection. The suggested *optTree* [KH□3] heuristics take the constrained multi-core scheme one step further by bringing in additional metric for the refinement of ultimate protocol performance.

□n interesting result as part of [SBK□□]'s studiesis that the tree rooted at core selected in consideration of the source group incurred less cost than the tree rooted at the core selected close to center of the receiver set, with the gain in tree cost increasing as the number of sources increase. This result validates the fact that a multicast tree rooted at a single core to optimi□e the paths to the receivers rather than those from the sources provides more gain in efficiency for optimi□ing the shared paths to the receivers rather than the unique paths from each source to the core. Note that this result is less significant for multi-core architectures which allow the placement of cores according to the proximity of multicast group members partitions in the domain, thereby optimi□ing the incurred cost locally on the tree portion covering the member partition.

Calvert et al [CZD□□]'s topological center results are significant for single-core selection in homogeneously distributed applications for separating the selection process from the multicast group information and allowing the use of a core selected statically in the domain. Topological center contributes especially for the case the frequency of changes in group membership is high in a well-distributed group, reducing the necessity for the core migration to the instance of core failures.

The candidate core set processed by each algorithm is indicative for the message exchange and processing overhead during the selection process. In certain algorithms the selection method inherently restricts the candidate core set to a particular set of nodes in the domain. *Tournament* [SBK□□] and*INITIAL* [Sa□6], examining only the nodes which are in the subgraph □bordered□ by the receiver group are in this group. Group-based core choices [CZD□□] restrict the candidate core set to the set of receivers. Candidate core set of a smaller si□e can also be imposed

externally to the ultimate selection module by restricting the candidate domain through a less excessive and less dynamic process to reduce the overhead still keeping the coverage of the set well-distributed throughout the network.

Core selection is a trade-off between the amount of processing and message transmission during the selection process, and the performance of the protocol in significant metrics. The performance results improved as the group size increased validating that as the size increases, the receivers are well-distributed over the network making the locations of the core less significant [TR97, CZD95]. According to Calvert et al [CZD95], group information does not necessarily improve performance. When the multicast group is widely distributed, the group-based core choice can degrade the performance relative to the topology-based methods. A certain core selection mechanism demonstrates different results for various applications. However a core selection method effective for one application is very likely to be effective in a similar manner for other types of applications compared to other selection mechanisms. There is no single-best core choice method across all metrics.

5. Conclusions

In this paper, we reviewed and analyzed the core-based approach for core selection algorithms in multicast routing protocols. We started with a brief overview of multicast routing protocols and the current trend into core-based approach. We reviewed the core selection algorithms studied in literature for their algorithmic structures, application issues and performance results as tested by their respective authors. Our review resulted in classification of core-selection schemes for their operational characteristics. We then analyzed the computational and message complexity of each algorithm in availability of local distance-vector state information at domain nodes. Our discussion included remarks on the applicability of these algorithms into the available multicast routing protocols.

Multicast routing protocols are mainly designed for their feasible deployment in distributed domains. Development and maintenance of the multicast tree structure above a certain level in the measures described in its design goal is determinant on a protocol performance. The existing architectures can be enriched with components for selection and management of core set for improved approximation of the generated tree and thus protocol operations. In this paper, we examined core selection, placement of core(s) to govern the multicast path for its placement in domain, and its construction especially in multi-core case. Other areas of similar concern are measurement of the tree deterioration from the optimal or approximated structure over the course of multicast session, and core migration for keeping the tree quality above a certain level according with its design description.

The design goals fulfilled by the core-based algorithms are governed by those of the routing protocols, and the assumed design problem in a multicast routing protocol existing in literature is the development of shortest paths between the tree root(s), namely the core(s) in core-based architecture, and the receivers. Multicast routing protocols and most core selection algorithms assume and operate on a unique metric. However, today's QoS-demanding applications require consideration of multiple metrics in optimization, and the typical design problem is the