Impact of Simple Cheating in Application-Level Multicast [1]

INFOCOM 2004

presented by wdchiang

Outline

- Introduction
- Application Layer Multicast
  - Protocol : HBM, TBCP, NICE, NARADA
  - Cheating
- Effects of Cheating in Application Layer Multicast
  - Network Stress Ratios, Stretch Ratios
- Discussions
- References

Introduction

- Most of the existing application layer multicast protocols require that each node measures its distance measurements to other nodes and reports these measurements for decision making.
- A cheat can therefore lie about its distance measurements, in order to be positioned closer to the data source while limiting the replication burden.

Introduction

- Although the problems of cheats has often been mentioned, we are not aware of any quantitative study of their effects.
- The paper seeks to extract possible common consequences and trends created by the presence of cheats in application layer multicast overlay trees.

Application Layer Protocol

Protocols	Algorithm	Distributed Algorithm	Scope
HBM [2]	Centralized		Full
TBCP [3]	Distributed	Tree-First	Limited
NICE [4]	Distributed	Tree-First	Full
NARADA [5]	Distributed	Mesh-First	Full

Application Layer Protocol: HBM

- Principles
  - The construction and maintenance of the overlay tree is under the control of a single host, RP.
  - Periodically and asynchronously, each group member measures its distance to all the others and reports to the RP.
  - The RP is then responsible for the overlay topology calculation and dissemination among the group members.
  - The topology used in this study is a degree-bounded shared tree of minimum cost, based on RTT.

Application Layer Protocol: HBM

- Principles
  - Redundant Virtual Links

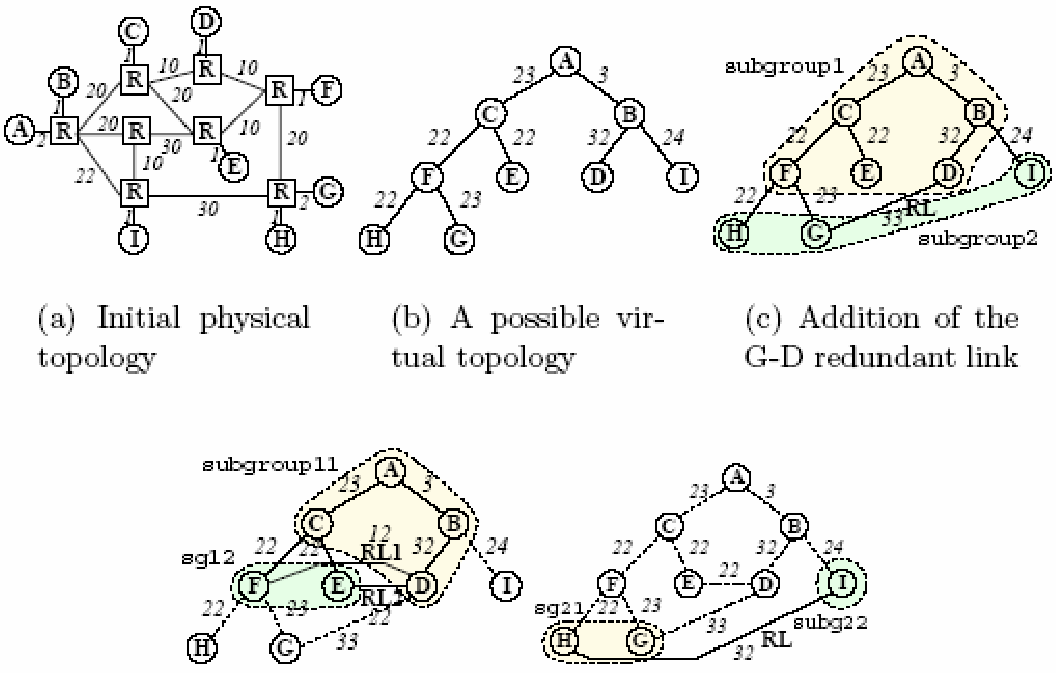

(a) Initial physical topology

(b) A possible virtual topology

(c) Addition of the G-D redundant link

(d) Addition of E-D redundant link

(e) Addition of the H-I redundant link

Fig. 4. Addition of Redundant Virtual Links (RVL), an example.

Application Layer Protocol: HBM

- Simple Cheating Method
  - An HBM cheat always reports a distance of 0 to the source, and adds 10 seconds to the RTT distances it measured to the rest of the group.
  - An HBM cheat also delays by 10 seconds any measurement probes it receives from any other group member.
  - A cheat is thus aiming to become one of the source’s children, while having less children.

Application Layer Protocol: TBCP

- Principles
  - Join Procedure

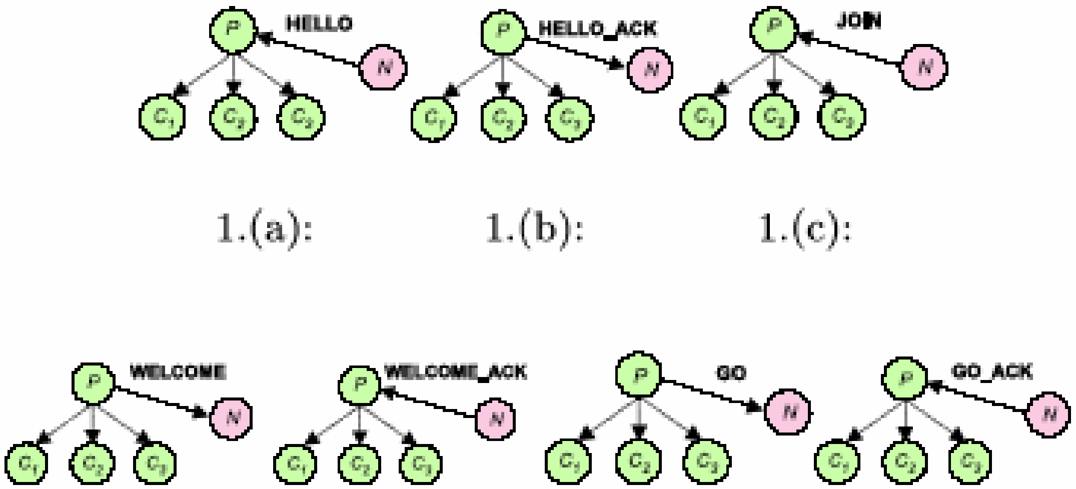

1.(a): 1.(b): 1.(c):

1.(d): 1.(e): 1.(f): 1.(g):

Fig. 1. TBCP Join procedure messages

Application Layer Protocol: TBCP

- Principles
  - Local Configuration

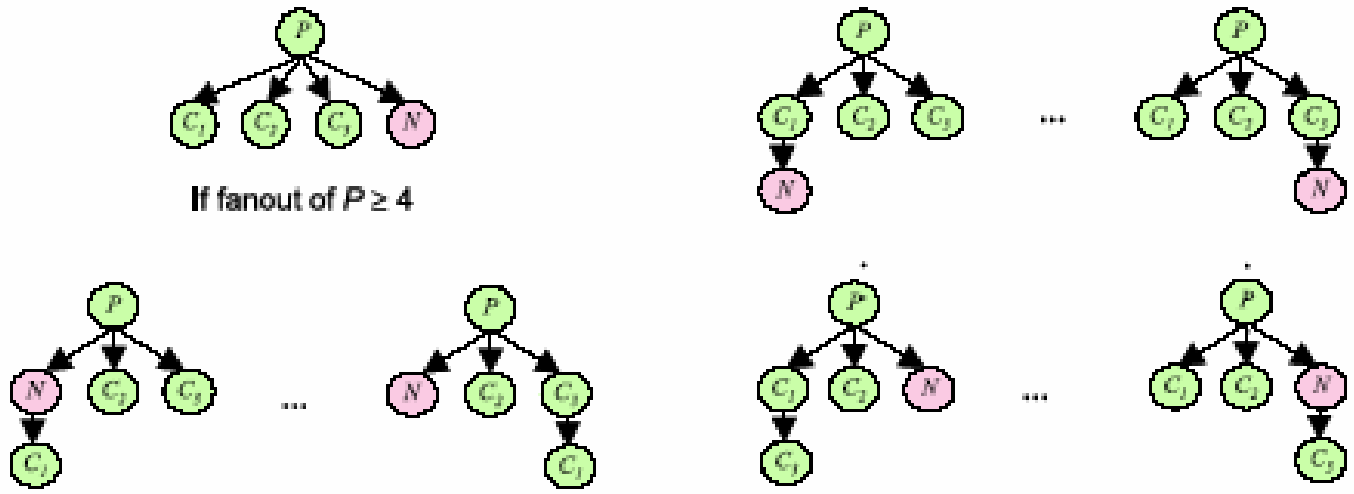

Fig. 2. Local configurations tested.

Application Layer Protocol: TBCP

- Principles
  - It is a recursive algorithm where, starting at the tree root as a potential parent, a newcomer measures the distance between it self and potential parent, along with the distance between itself and all its potential siblings.
  - The distances are reported to the potential parent.
  - The potential parent then considers all the local configurations for the acceptance of the newcomer in the tree.

Application Layer Protocol: TBCP

- Principles
  - The best local configuration is chosen and the appropriate node directed to its “next” potential parent where the algorithm starts again.
  - TBCP has a maintenance method where nodes periodically “re-join” one of its known ancestors chosen at random.

Application Layer Protocol: TBCP

- Simple Cheating Method
  - A TBCP cheat will always report a distance of zero to its potential parent.
  - A cheat will choose the minimum allowed fan-out value.
  - However, to try and avoid having a child, cheats also lie about their distance to other receivers: a cheat always delays a received probe by a fixed amount of time(10 sec.) and always add a fixed amount of time to the distance it reports from other receivers.

Application Layer Protocol: NICE

- Principles
  - In NICE, nodes arrange themselves into a hierarchy of clusters whereby clusters belong to layers and nodes belonging to a cluster are close to each other in relation to some given cost metric.
  - All nodes belong to a cluster in the lowest layer of the hierarchy but cluster leaders are also members of a cluster in their next-higher layer.

Application Layer Protocol: NICE

- Principles

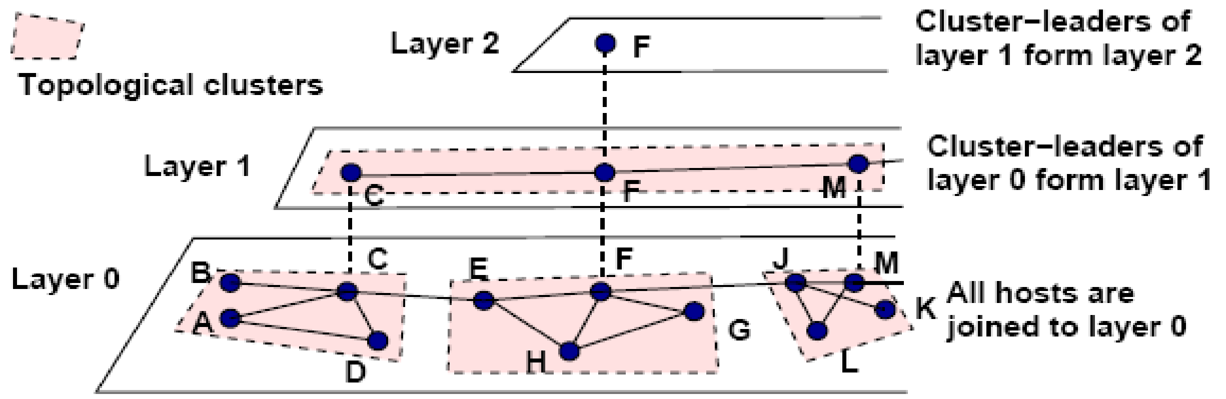

Fig. 2. Hierarchical arrangement of hosts in NICE. The layers are logical entities overlaid on the same underlying physical network.

Application Layer Protocol: NICE

- Principles
  - Join Procedure

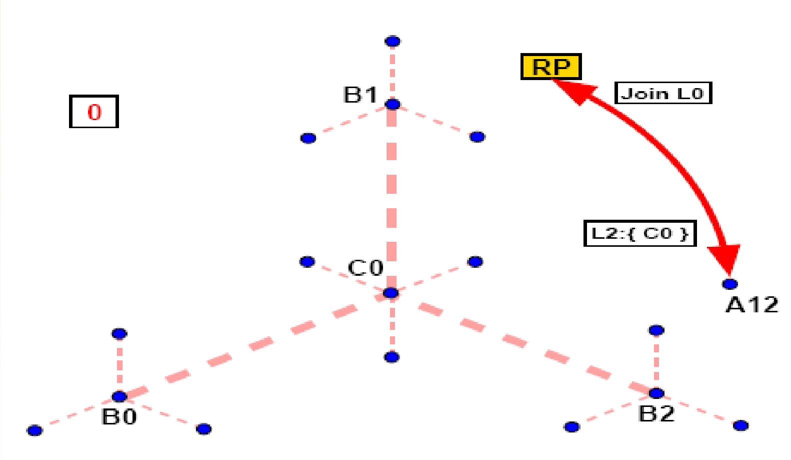

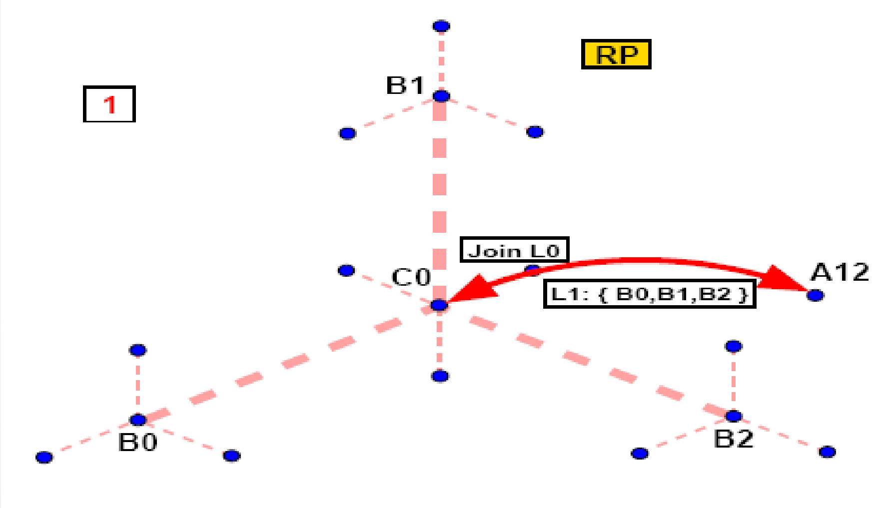

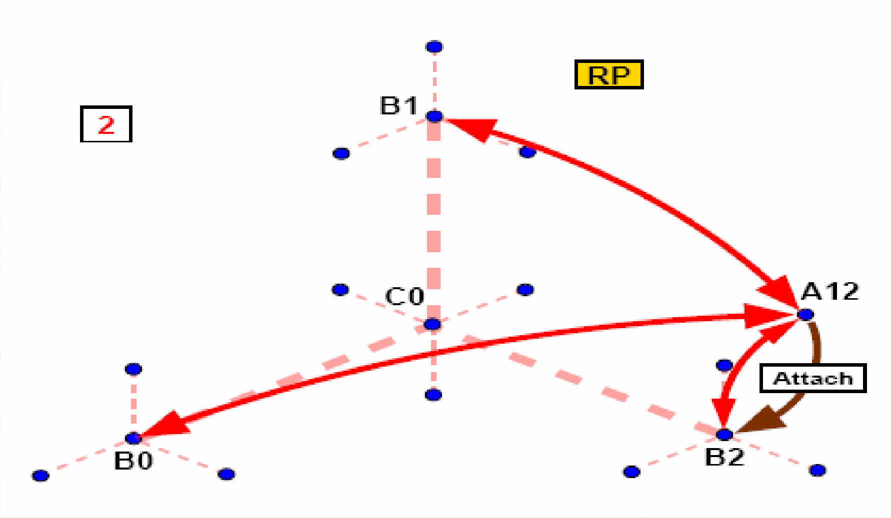

Application Layer Protocol: NICE

- Principles
  - Members of a cluster periodically exchange heartbeat messages with each other containing an estimate of the distance from themselves to each of the other cluster members.
  - Whenever membership of a cluster changes the cluster leader, using this cluster member distance information, checks if it is still the center of the cluster and thus the most appropriate leader, transferring leadership to another cluster member if necessary.

Application Layer Protocol: NICE

- Principles
  - The cluster leader periodically checks the size of its cluster and splits the cluster if its membership exceeds an upper bound.
  - If the cluster size falls below a lower bound, the leader merges its cluster with the closest cluster belonging to the next-higher layer.

Application Layer Protocol: NICE

- Simple Cheating Method
  - A NICE cheat sets out to join a cluster in the highest layer possible to minimize its distance to the data source.
  - For a cheat to join a cluster in the next-higher layer it must become the leader of its highest-layer cluster and so tries to achieve leadership through quoting, in its heartbeat messages, only a fraction of the actual distance to the other cluster members.
  - On recalculating which node is closest to all the others, the current cluster header will likely transfer its leadership to the cheat.

Application Layer Protocol: NICE

- Simple Cheating Method
  - Once a cheat has gained leadership of a cluster it will make sure never to transfer leadership from itself to any other cluster members, by reporting a distance of 0 to all other cluster members in heartbeat messages.
  - A cheat will never merge its clusters in the lower layers and will also delay cluster join requests from other nodes by 10 seconds to reduce the likelihood of these joining the clusters.

Application Layer Protocol: NARADA

- Principles

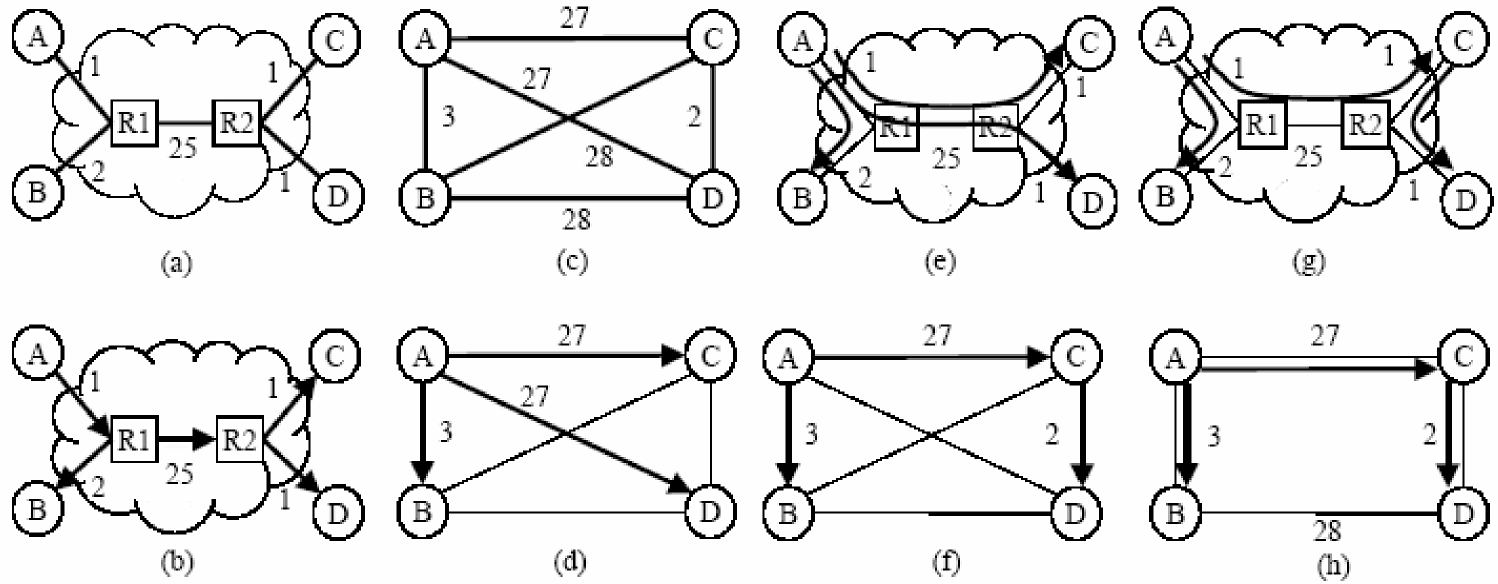

Figure 1: Examples to illustrate IP Multicast, naive unicast and End System Multicast

Application Layer Protocol: NARADA

- Simple Cheating Method
  - A cheat will set out to add the source as a mesh-neighbor and so receive data directly from the source.
  - A cheat makes sure to establish at least one mesh link to another node through which it misleads the source to believe that it can deliver data to all the other at a fraction of the actual costs.
  - NARADA is susceptible to partitioning when the degree of mesh nodes is small, so in order not to break the protocol and no more.

Effects of Cheating in Application Layer Multicast

- The Stress
  - The Stress is defined as the number of redundant copies of a data packet carried on a network link.
  - The maximum stress is the maximum number of duplicates seen by any single network link.
  - The average stress is the sum of duplicates divided by the total number of network links.
  - A major goal of all protocols is to keep the value of these stress indicators as small as possible.

Effects of Cheating in Application Layer Multicast

- Network Stress Ratio
  - stress ratio = stress/stress$_{ref}$, where stress$_{ref}$ is the corresponding stress observed when all receivers behave in an honest way.
  - Note that a stress ratio smaller (resp. greater) than 1 represents an improvement (resp. deterioration) compared with the case without any cheat.

Effects of Cheating in Application Layer Multicast

- Network Stress Ratio

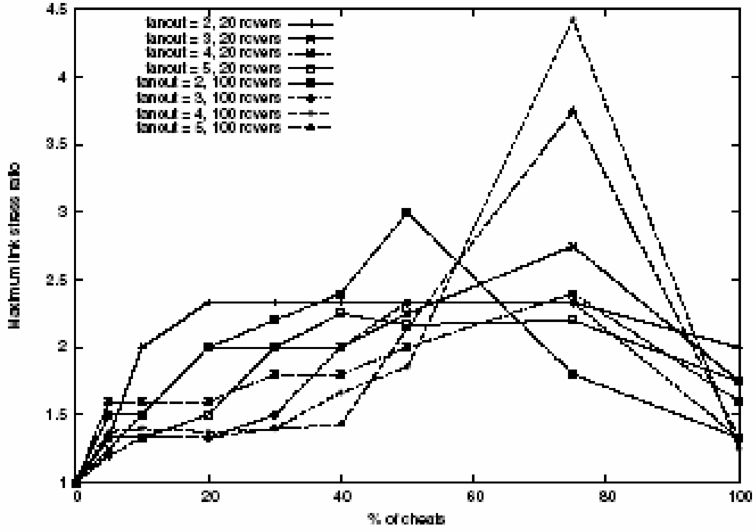

Fig. 1. Maximum link stress ratios in HBM

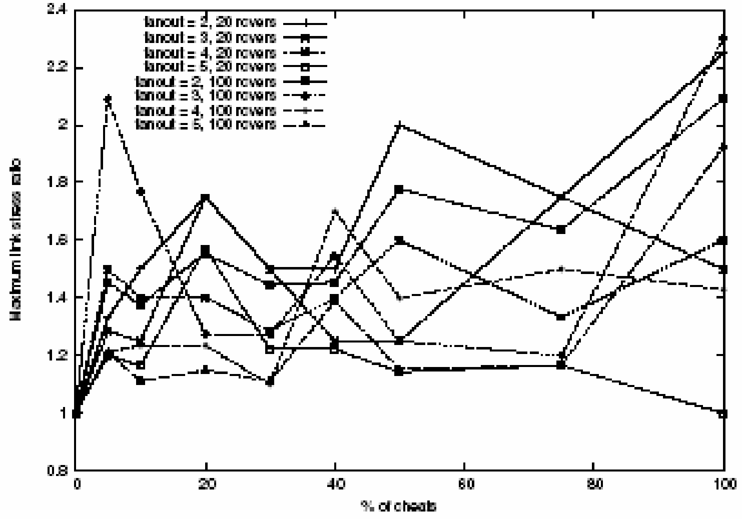

Fig. 2. Maximum link stress ratios in TBCP

Effects of Cheating in Application Layer Multicast

- Network Stress Ratio

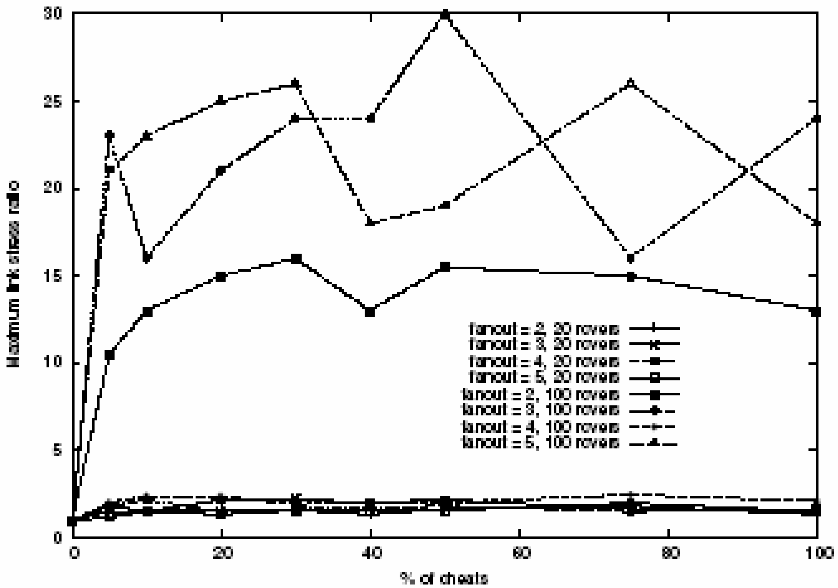

Fig. 3. Maximum link stress ratios in NICE

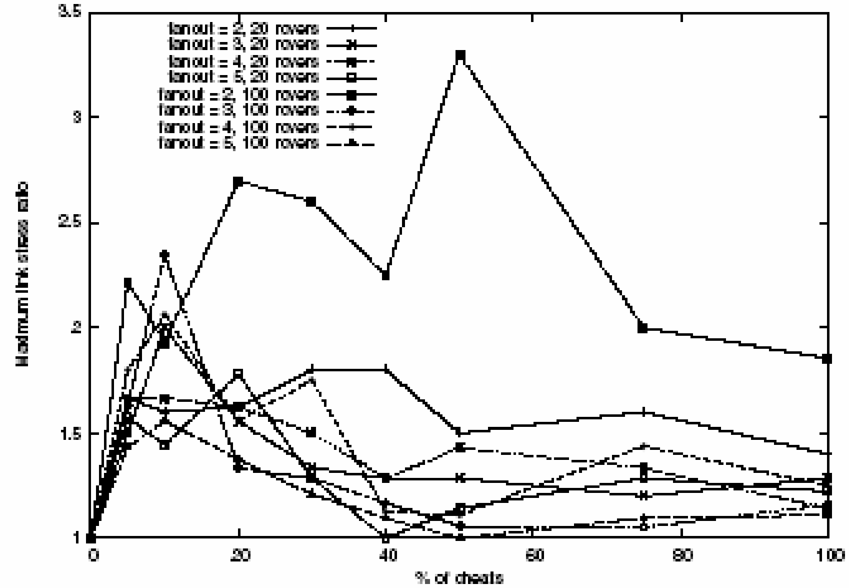

Fig. 4. Maximum link stress ratios in NARADA

Effects of Cheating in Application Layer Multicast

- The Stretch
  - The stretch, or relate delay penalty (RDP), is a measure of the penalty paid by a receiver for receiving data on an application layer tree rather than directly from the source.
  - It is defined as the ratio TD/UD, where TD is the tree delay and UD is the unicast delay.
  - stretch ratio = $stretch/stretch_{ref}$, where $stretch_{ref}$ is the stretch of a receiver observed when all receiver behave in an honest way.

Effects of Cheating in Application Layer Multicast

- Stretch Ratio
  - the average stretch ratio for honest receivers in the presence of cheats

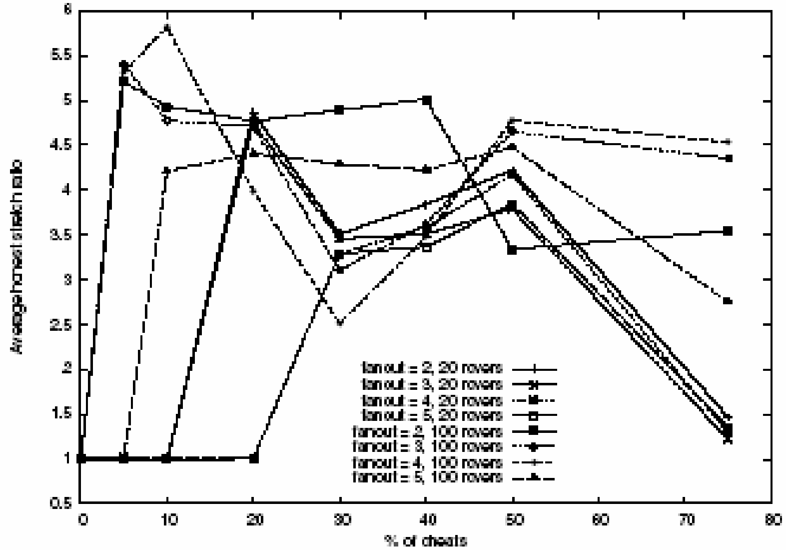

Fig. 5. Average stretch ratios for honest receivers in HBM

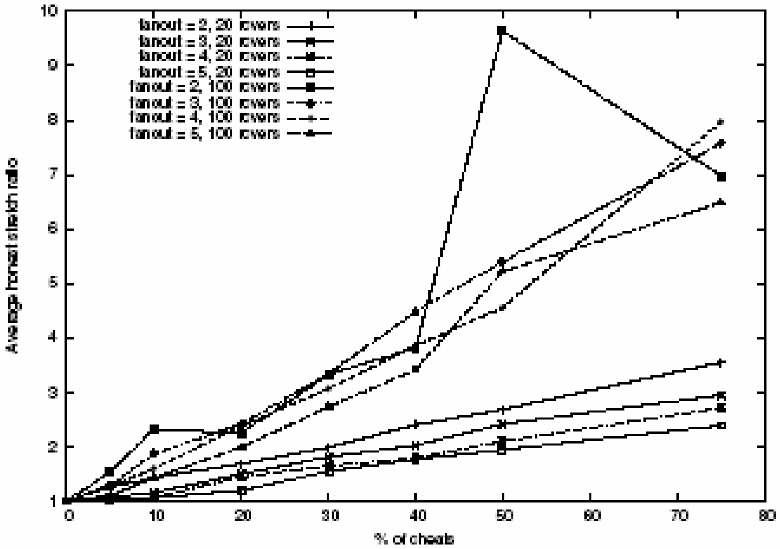

Fig. 6. Average stretch ratios for honest receivers in TBCP

Effects of Cheating in Application Layer Multicast

- Stretch Ratio
  - the average stretch ratio for honest receivers in the presence of cheats

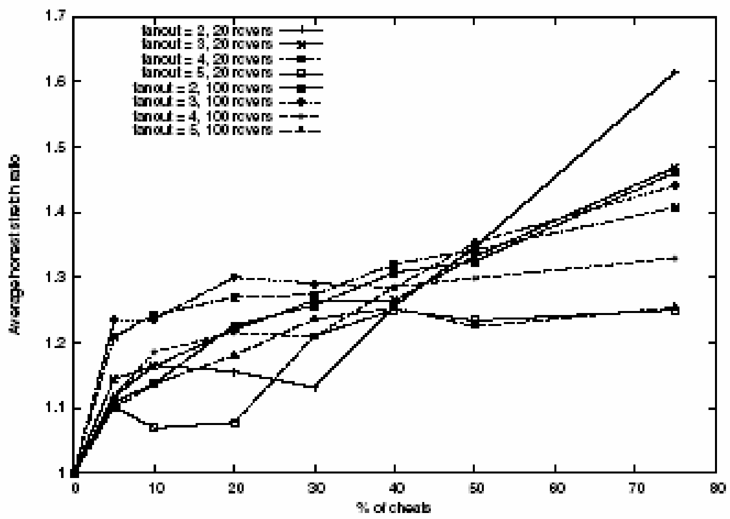

Fig. 7. Average stretch ratios for honest receivers in NICE

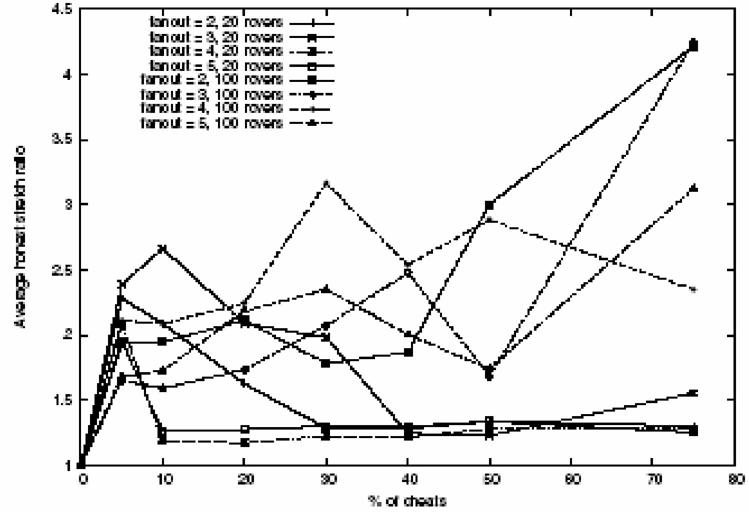

Fig. 8. Average stretch ratios for honest receivers in NARADA

Effects of Cheating in Application Layer Multicast

- Stretch Ratio
  - the average ratio in stretch for the cheats themselves

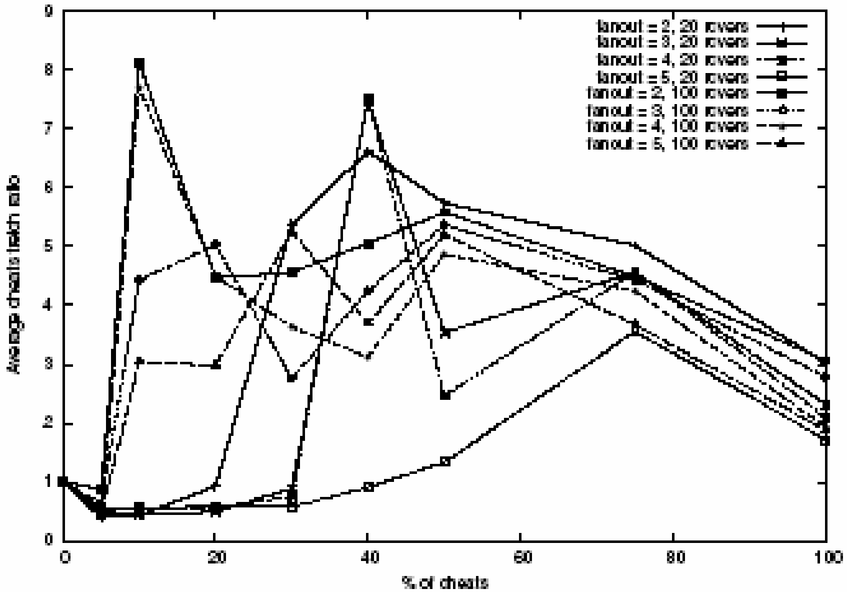

Fig. 9. Average stretch ratios for cheats in HBM

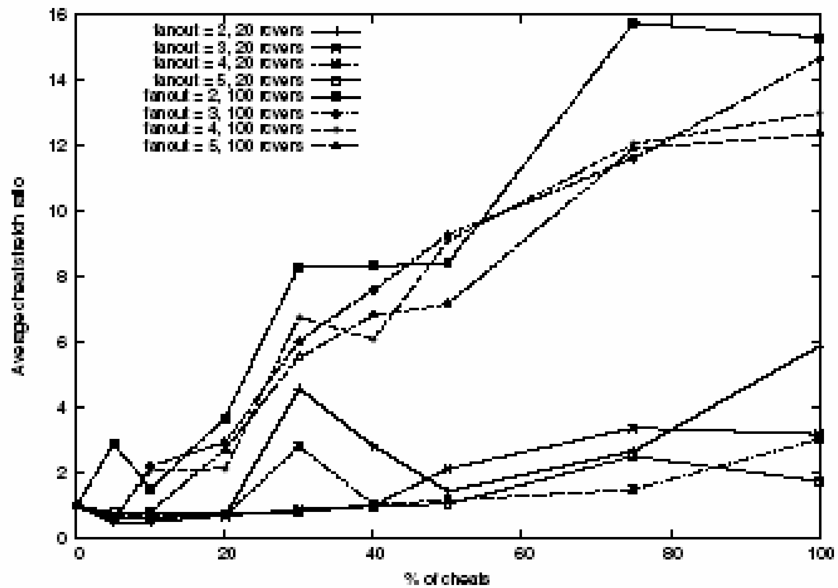

Fig. 10. Average stretch ratios for cheats in TBCP

Effects of Cheating in Application Layer Multicast

- Stretch Ratio
  - the average ratio in stretch for the cheats themselves

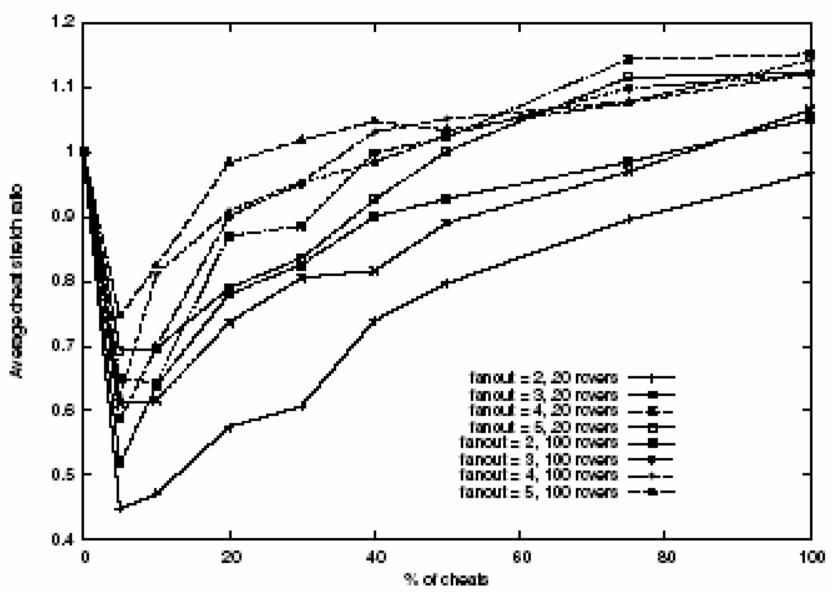

Fig. 11. Average stretch ratios for cheats in NICE

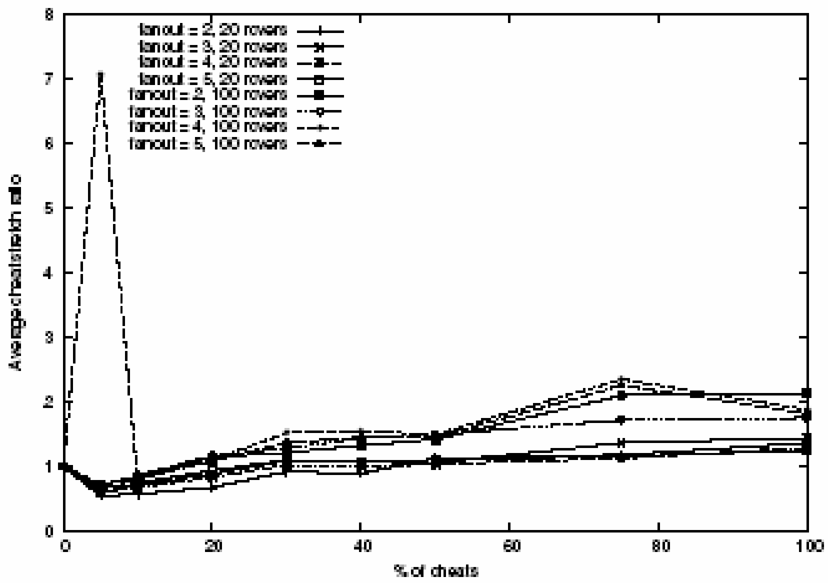

Fig. 12. Average stretch ratios for cheats in NARADA

Effects of Cheating in Application Layer Multicast

- Cheats vs Random Tree

group size	fanout	% of cheats
20	2	20 – 30
20	3	30 – 40
20	4	30 – 40
20	5	40 – 50
100	2	5 – 10
100	3	5 – 10
100	4	5 – 10
100	5	5 – 10

TABLE I

HBM CHEATS VS RANDOM TREE: % OF CHEATS WHEN CHEATS ARE BETTER OFF BEING IN RANDOM TREE.

group size	fanout	% of cheats
20	2	never
20	3	75 – 100
20	4	40 – 50
20	5	40 – 50
100	2	75 – 100
100	3	40 – 50
100	4	30 – 40
100	5	20 – 30

TABLE III

NICE CHEATS VS RANDOM TREE: % OF CHEATS WHEN CHEATS ARE BETTER OFF BEING IN RANDOM TREE.

group size	fanout	% of cheats
20	2	20 – 30
20	3	30 – 40
20	4	20 – 30
20	5	40 – 50
100	2	0 – 5
100	3	5 – 10
100	4	5 – 10
100	5	10 – 20

TABLE II

TBCP CHEATS VS RANDOM TREE: % OF CHEATS WHEN CHEATS ARE BETTER OFF BEING IN RANDOM TREE.

group size	fanout	% of cheats
20	2	40 – 50
20	3	20 – 30
20	4	20 – 30
20	5	20 – 30
100	2	10 – 20
100	3	10 – 20
100	4	0 – 5
100	5	10 – 20

TABLE IV

NARADA CHEATS VS RANDOM TREE: % OF CHEATS WHEN CHEATS ARE BETTER OFF BEING IN RANDOM TREE.

Discussions

- Actually, none of these protocols were explicitly designed o deal with cheats, and all showed, at various point of the study, that their performance could quickly degrade to be worse.
- Designing general cheat detection and prevention techniques for various types of metrics is an research challenge.

References

[1] Mathy L., Blundell N., Roca V., El-Sayed A., Impact of Simple Cheating in Application- Level Multicast. In INFOCOM 2004

[2] V. Roca and A. El-Sayed. A Host-Based Multicast Solution for Group Communications. In IEEE Intl. Conf. Networking, Jul 2001.

[3] L. Mathy, R. Canonico, and D. Hutchison. An Overlay Tree Building Control Protocol. In Proc. of Intl. workshop on Networked Group Communication (NGC), pages 76–87, Nov 2001.

[4] S. Banerjee, B. Bhattacharjee, and C. Kommareddy. Scalable Application Layer Multicast. In ACM SIGCOMM, Aug 2002.

[5] Y-H. Chu, S. Rao, and H. Zhang. A Case for End System Multicast. In ACM SIGMETRICS, pages 1–12, Santa Clare, CA, USA, June 2000.