Dynamic Characteristics of *k-ary n-cube* Networks for Real-time Communication

Gerald Fry and Richard West
Boston University
Boston, MA 02215
{gfry,richwest}@cs.bu.edu

Computer Science

Introduction

- Overlay topologies have become popular in Peer-to-peer (P2P) systems
  - Efficiently locate & retrieve data (e.g., mp3s)
  - e.g., Gnutella, Freenet, Kazaa, Chord, CAN, Pastry
- Previous work
  - *Static* analysis of *k-ary n-cube* graphs for structuring overlay topologies
  - Basis stems from interconnection networks in parallel architectures, such as SGI Origin 2/3000
  - Focus on delivery of real-time data streams

Dynamic Characteristics

Computer Science

- Internet-scale overlays are inherently more dynamic than tightly coupled interconnection networks:
  - Transient mappings of physical host addresses to logical node identifiers
  - Adaptation of logical positions of publishers and subscribers to satisfy QoS constraints
  - Hosts joining and departing from the system
- Must adapt the overlay structure to maintain connectivity and optimal hop count

Contributions

- Focus on *dynamic* analysis of *k-ary n-cube* logical networks
  - Methods for performing *M-region transitions*
  - Calculation of message exchange overheads for join and departure events
  - Quantify lag effects of join/departure *bursts*
- Applications: live video broadcasts, resource intensive sensor streams, data intensive scientific applications

Properties of k-ary n-cube Graphs

Computer Science

- $M = k^n$ nodes in the graph
- If k = 2, degree of each node is n
- If k > 2, degree of each node is 2n
- Worst-case hop count between nodes:
  - $n\lfloor k/2 \rfloor$
- Average case path length:
  - $A(k,n) = n \lfloor (k^2/4) \rfloor 1/k$
- Optimal dimensionality:
  - n = ln M
  - Minimizes A(k,n) for given k and n

M-region Analysis

- Hosts joining / leaving system change value of the number of physical hosts, m
  - Initial system is bootstrapped with overlay that optimizes A(k,n)
- Let M-region be range of values for m for which A(k,n) is minimized
  - As m changes, *M-region transitions* restructure the overlay to maintain optimality
- In previous work, we derive the first sixteen M-regions

Join/Departure Events

- Each join/departure event requires a number of control messages to be exchanged
  - Necessary for maintaining global consistency of routing state
- An individual event requires O(n) message exchanges
  - Worst case message exchange overhead is independent of the occurrence of an *M-region transition*

Adaptation Lag

- Assume join/departure events arrive in bursts of size b
- Given the dimensionality, n, of the overlay, the *adaptation lag* = n • b • C
  - C is a constant proportional to the average transmission time of a single message
- Simulation investigates effects of adaptation lag due to join/departure burst requests
  - MMPP used to generate bursts
  - Target and actual M-regions recorded at discrete time intervals

Simulation Results

Computer Science

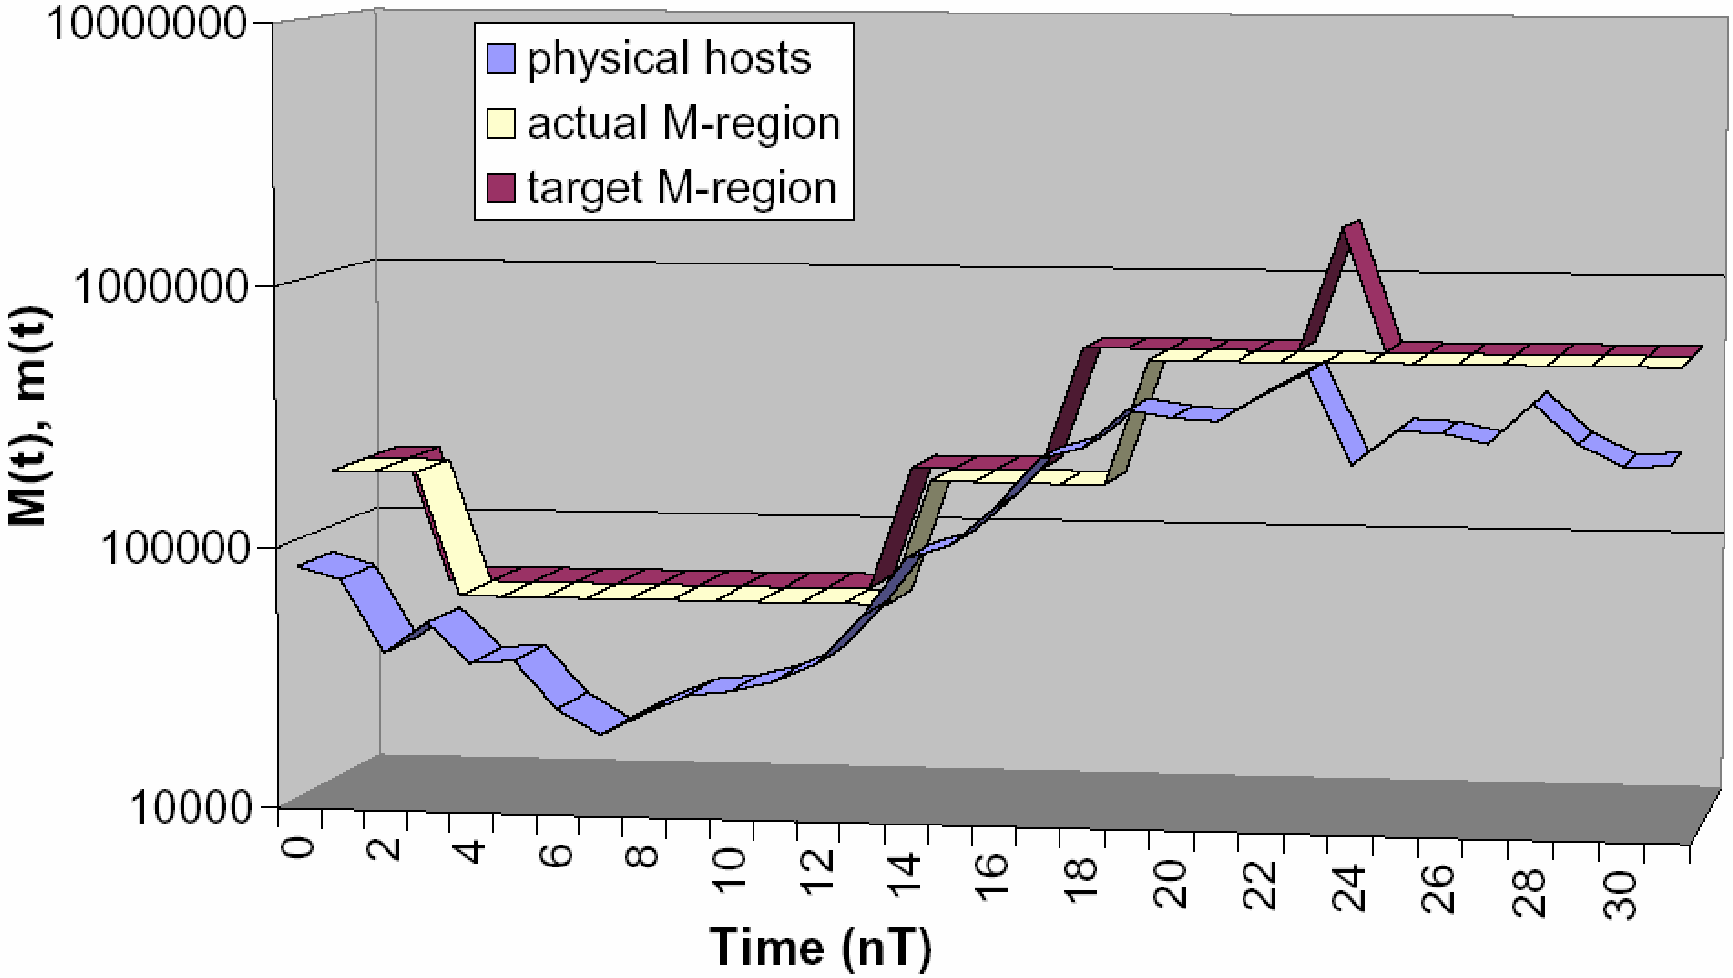

C = 10^{-6}, Target M-region utilization = 84.4%

Conclusions/Future Work

Computer Science

- Analysis of dynamic characteristics of *k-ary n-cube* overlays
  - Routing state maintenance due to host joins and departures
  - Optimal topology wrt. average and worst-case hop count as the system evolves
- Future Work
  - Build scalable overlay topologies and analyze performance in practice
  - Integrate end-host architectures for user-level sandboxing
  - Multicast tree construction in *k-ary n-cube* networks