MULTI-SCALE INTEGRATED INFORMATION AND TELECOMMUNICATIONS SYSTEM (MIITS)

FIRST RESULTS FROM A LARGE-SCALE END-TO-END NETWORK SIMULATOR

Lukas Kroc
Cornell University
Stephan Eidenbenz, James P. Smith, Roman Waupotitsch
Los Alamos National Laboratory

Los Alamos NATIONAL LABORATORY EST. 1943
Operated by the Los Alamos National Security, LLC for the DOE/NNSA

Overview

- **Introduction**
  - Network simulation motivation
  - Approaches and challenges
- **MIITS**
  - System Overview
  - Performance considerations
- **Case study: Los Angeles IP Network**

Introduction: Motivation

- Communication networks:
  - Complex systems with *heterogeneous* components
  - Highly *inter-dependent* services (WWW, VoIP, P2P,...)
- Simulation:
  - Helps understand behavior of networks
  - Aims at mimicking key features of reality
- Ingredients:
  - Network topology and physical capabilities
  - Decision processes (protocols)
  - Communication requests (sessions)

Introduction: Approaches and Challenges

- Network topology
  - Known, but proprietary
  - Approximate: using tracerouting
- Protocols
  - Known, but very complex
  - Approximate: flows, ..., simplifications
- Sessions
  - Unknown, but can collect statistics
  - Model: statistically equivalent to real world

Introduction

- Goal
  - simulation of nation-wide communication infrastructure
  - “meaningful” and “expressive” results, providing “useful” information
  - incorporating all ingredients of a realistic simulation with tunable quantity/quality trade-off

- Our system
  - Multi-scale Integrated Information and Telecommunications System (MIITS)

System Overview: Core Design Considerations

- Extensibility, ease of use
  - New protocols and processes easily added
- Expressive power
  - We do not know what new approaches will require
- Conceptual simplicity
  - So that, as it grows, it doesn't become a jungle
- Scalability
  - Need to be able to run efficiently on parallel architectures

⇒ Simulation Framework library

System Overview: Network Simulation

- End-to-end approach
  - Creating realistic network
  - Topology and realistic communication sessions between end users by individually modeling end user devices and behavior
- Integrated approach across various types of networks
  - Integration of packet-switched and circuit-switched networks
  - Requires representation of end devices that can access different types of networks

System Overview

- Data flow
  - Models realistic communication sessions between end users
  - Models realistic network topology
  - Models individual end-user devices
  - Leverages existing socio-technical simulation technology (TRANSIMS)

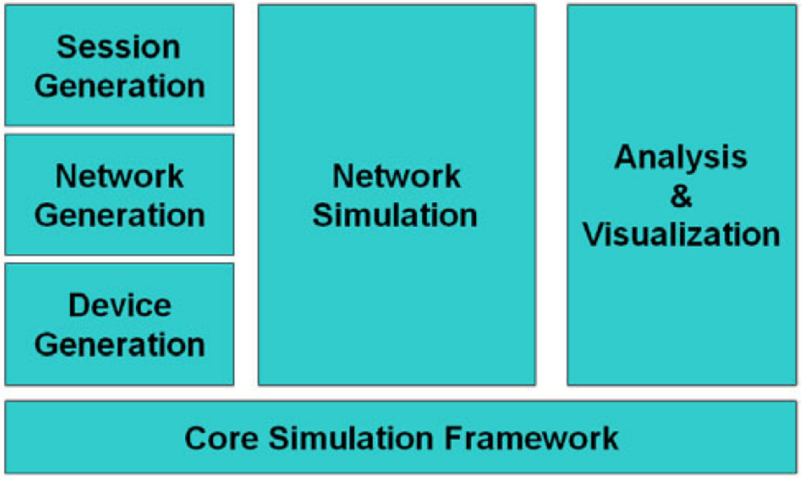

System Overview

- Network objects
  - Implements representations of *Devices*, *Interfaces*, and *Media*
  - Models every protocol layer of every device in the simulation
  - Different device types are created by combining different protocols on different layers of the protocol stack

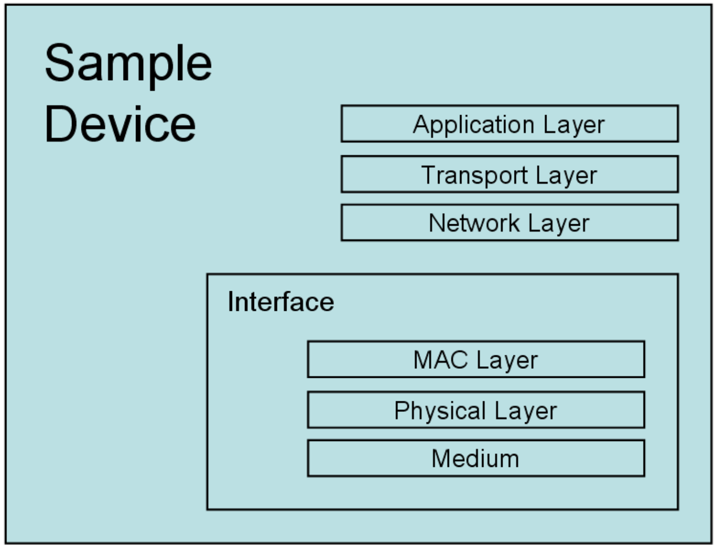

System Performance

- Memory requirements
  - ~3000 bytes per device, mostly stats and indexes
- Number of messages exceeds number of devices by 3 orders of magnitude
  - 1 million sessions result in 2.255 billion messages
  - both system control messages and IP packets
- Purely distributed implementation
  - no memory bottleneck
  - but excessive number of CPUs result in high communication overhead

System Performance

- Exhibits the distributed computation trade-off:
  - Few CPUs ⇒ little communication, too much computation
  - Many CPUs ⇒ little computation, too much communication

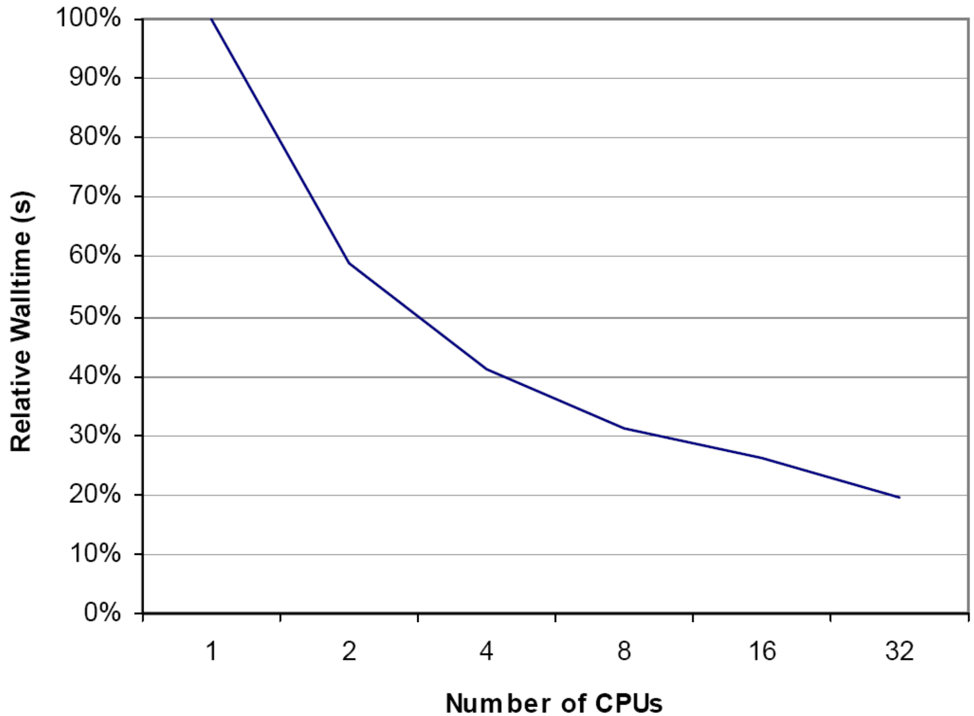

Los Angeles Network Simulation

- Motivation
  - Serves as a proof of concept of the MIITS system
  - Identify bottlenecks in the network
- Overview
  - IP network in Los Angeles area
  - Client/server communication pattern
  - Look at how the systems behaves with changing communication requirements
    - Uses some very simplified components

Los Angeles Network Simulation

- Network Topology:
  - Core network consists of 11225 nodes
    - Identified using *traceroute* from various national locations
  - ~1M end-user devices attached to the core network
    - some are clients, some servers

  $\Rightarrow$ Network represent ISPs in LA area, with attached users
- Sessions:
  - Average size 540kbyte per session
  - 1 million sessions
  - Average times between initiation of sessions: 1.5s, 0.5s, 0.1s, 10ms, 1ms, and 0.1ms

Los Angeles Network Simulation

- Results
  - Average packet travel time for session spacing
    - >= 10ms: 0.7 to 0.8 seconds (0.07 seconds per hop)
    - <= 1ms: up to 17 seconds
  - Reason is increased queue length in MAC layer
    - >= 10ms: max 11 packets
    - <= 1ms: max 512 packets

Los Angeles Network Simulation

- Results
  - Dropped packets

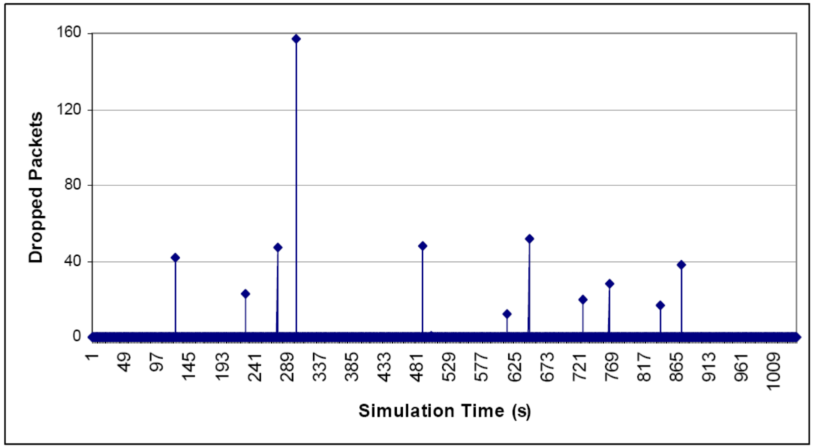

1ms session spacing

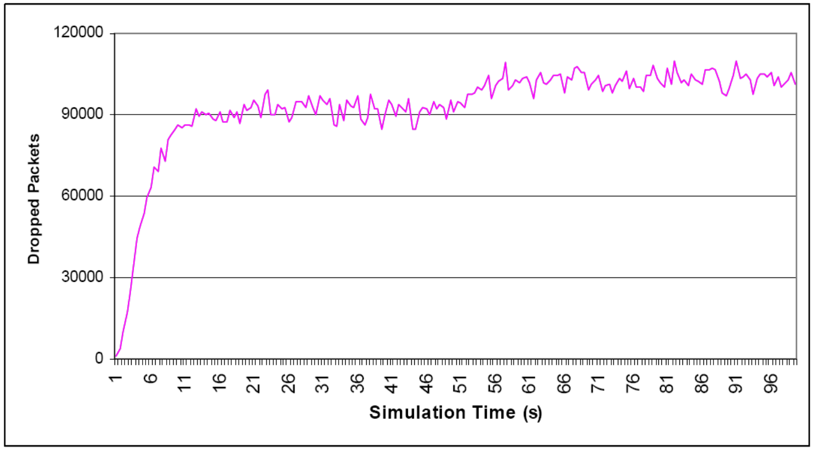

0.1ms session spacing

Conclusions

There is more to modeling communication systems than network simulation

- Steps towards integrated information system simulations
  - all ingredients are important: Network, Protocols, Sessions
- MIITS project:
  - Demonstrated prove of feasibility of large-scale packet-level simulation modeling the entire protocol stack
  - In the process of creating a national model
  - Work still in progress, mainly network and sessions

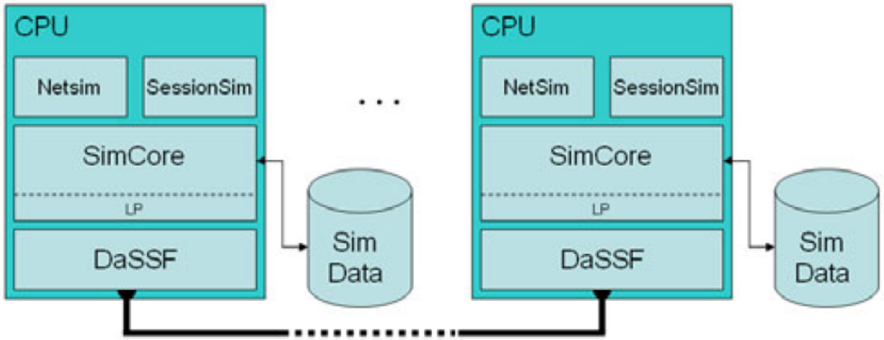
CPU
Netsim
SessionSim
SimCore
LP
DaSSF
Sim Data
CPU
NetSim
SessionSim
SimCore
LP
DaSSF
Sim Data

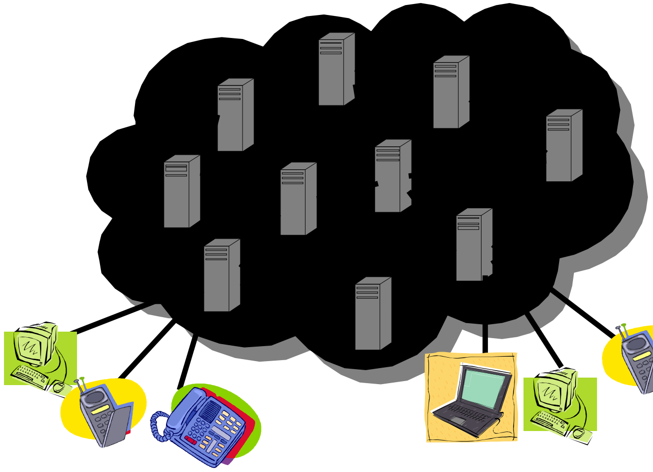

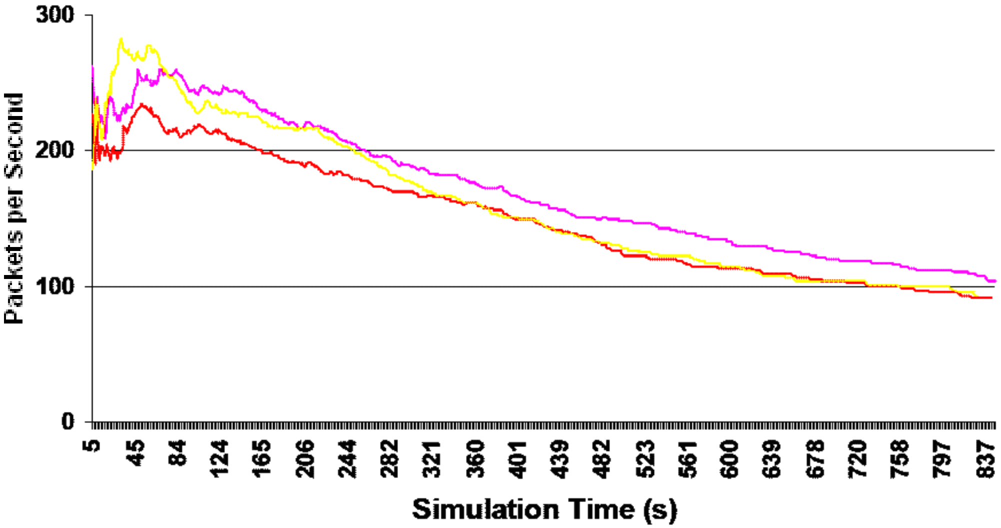
Packets per Second
300
200
100
0
5
45
84
124
165
206
244
282
321
360
401
439
482
523
561
600
639
678
720
758
797
837
Simulation Time (s)

Los Angeles Network Simulation

- Results
  - Queue lengths for session spacing of 1ms

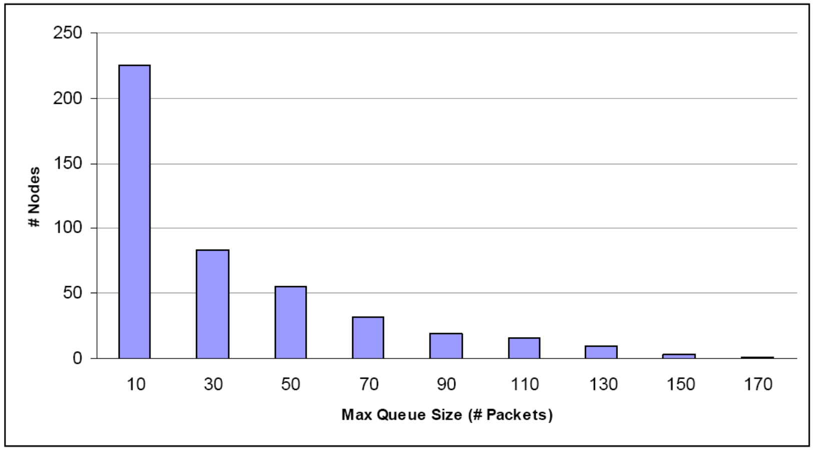

Distribution

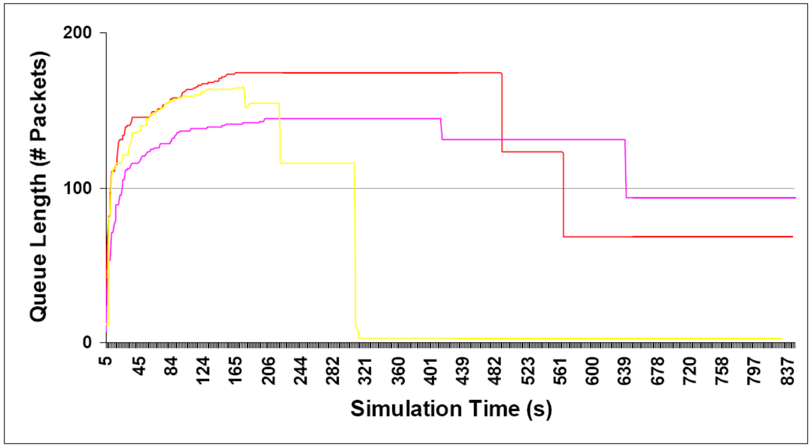

Maximum average load nodes

System Overview

- Scalability
  - Achieved through network aware implementation of all aspects of the system (except individual devices) to overcome memory and processing limitations
  - E.g. routing can be computed locally or distributed
  - Simulation core hides many of the complexities of of distributed systems

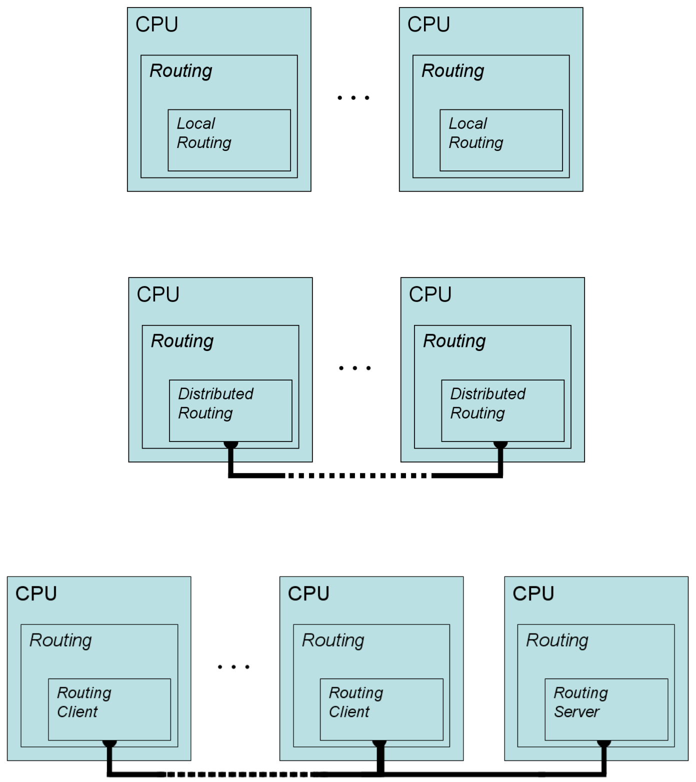

Los Angeles Network Simulation

- Scenario
  – Using subset of randomly identified clients and servers
  – Sessions:
    – UDP sessions, packets 60Kbits in size
    – 30% of sessions send 100 packets
    – Average is 72 packets or 540kbyte per session
    – 1 million sessions
  – Focus: Analyze network performance for different network loads
    – Average times between initiation of sessions: 1.5s, 0.5s, 0.1s, 10ms, 1ms, and 0.1ms

System Overview

- Integrated approach
  - Allows network simulations in different resolutions allowing the user to make tradeoff between accuracy and turn-around-time
    - E.g. Flow based simulation
    - E.g. Discrete Packet level simulation