Our Research on Real-

Raj Jain

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
http://www.cse.wustl.edu/~jain/

Overview

- Tasks:
  1. Multipoint Communication
  2. Virtual Source/Virtual Destination
  3. Real-Time ABR

Point-to-Point ABR

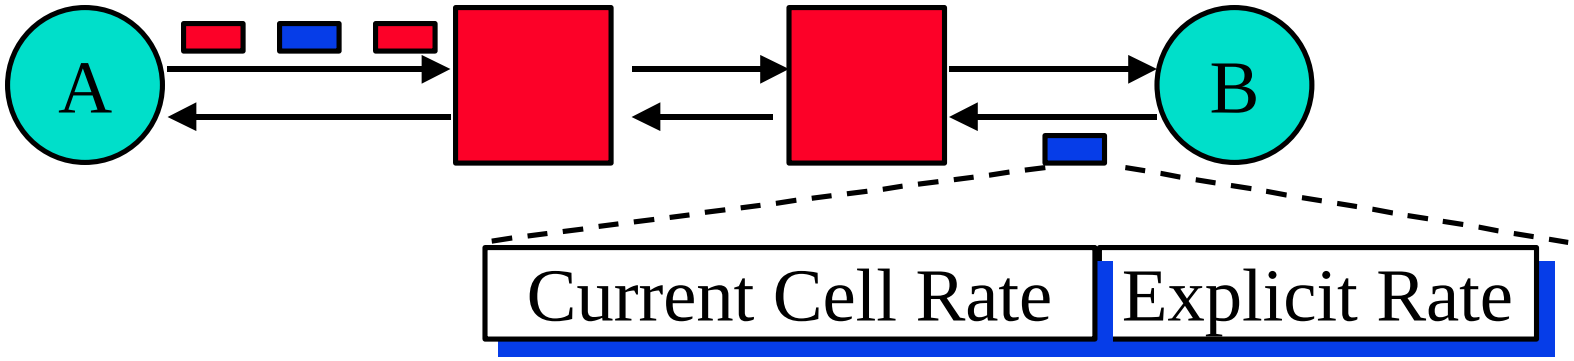

- Sources send one RM cell every n cells
- The RM cells contain "Explicit rate"
- Destination returns the RM cell to the source
- The switches adjust the rate down
- Source adjusts to the specified rate

Point-to-Multipoint ABR

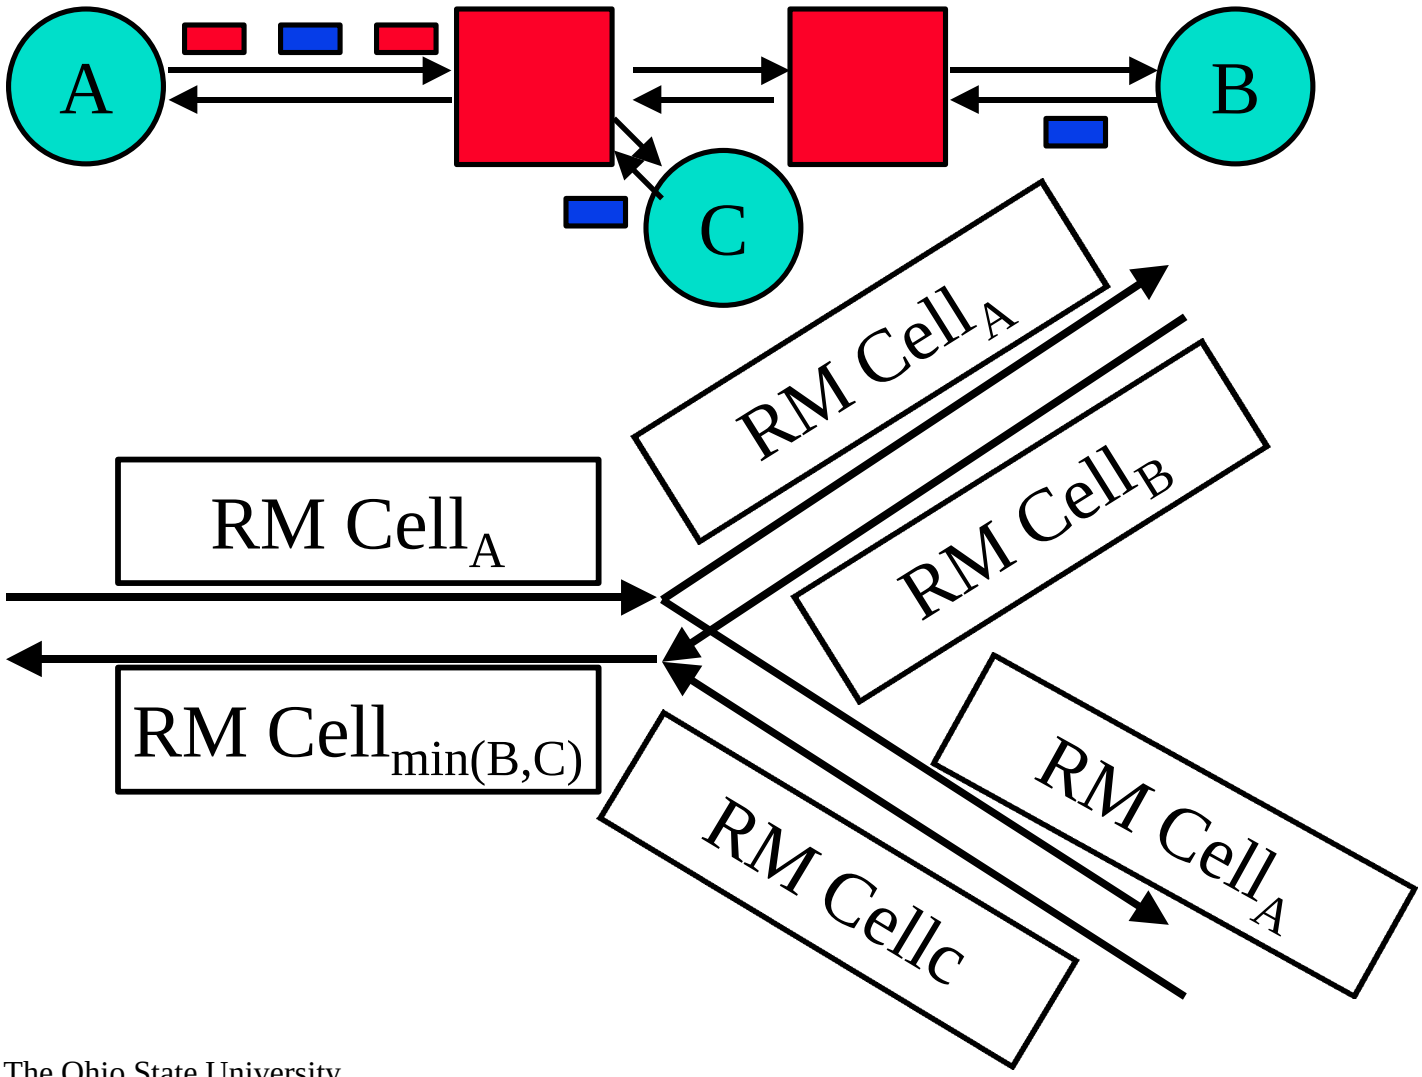

Point-to-Multipoint Connections: Issues

- Minimum of ER from branches is sent upstream. Should we wait for all branches?
- If you send BRM on every FRM, you may give feedback without receiving any
  ⇩ Need to ensure that at least one feedback has been received before sending a BRM.
  Otherwise, you may give PCR
- Not all downstream feedbacks in an upstream feedback ⇩ Consolidation noise

Basic Pt-Mpt: Results

- ABR with ERICA (extended for multipoint) works ok
- Efficiency, fairness, responsiveness is maintained
- Consolidation noise due to asynchronous arrival of feedback from different leaves appears as oscillations
- Additional delay due to FRM wait and BRM consolidation
  - ⇩ slower transient response than point-to-point
- Minimum of all paths is allocated
  - ⇩ Some links are underutilized
- Queue control (ERICA+) is required for stability

Impact I

- A summary of our ATM Forum contribution 97-0615 has been included in the base line for the next phase (TM4.1) of ATM Forum Traffic Management Specification
- Several leading industry members expressed interest in results

Multipoint-to-Point VCs

- A multipoint-to-point VC can have more than one concurrent sender
- Traffic at root = Σ Traffic originating from leaves

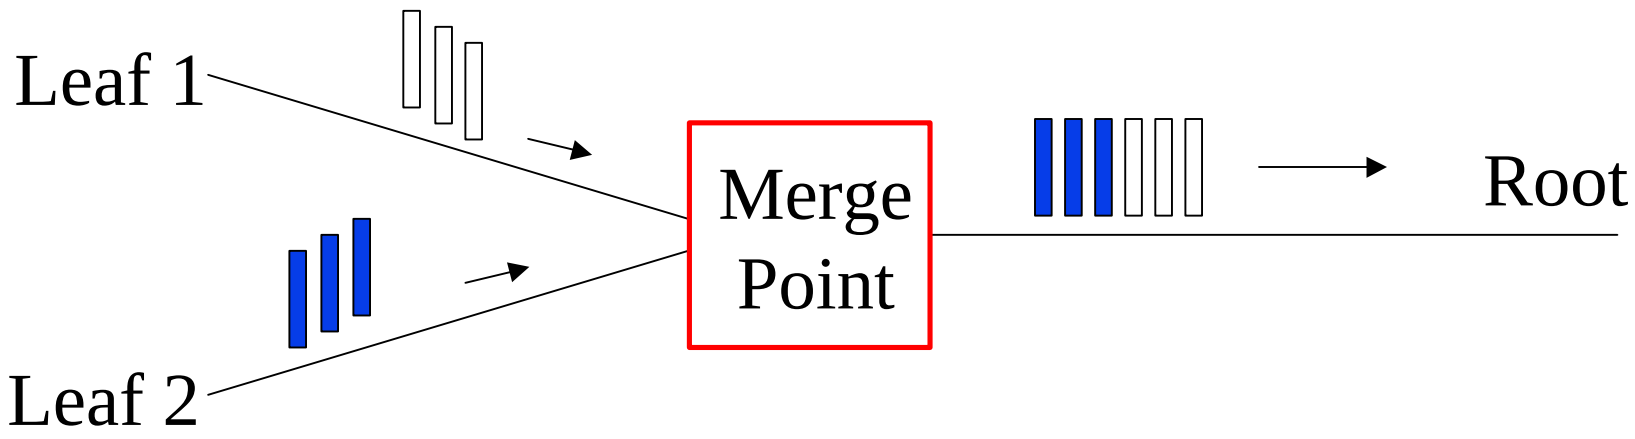

Sources, VCs, and Flows

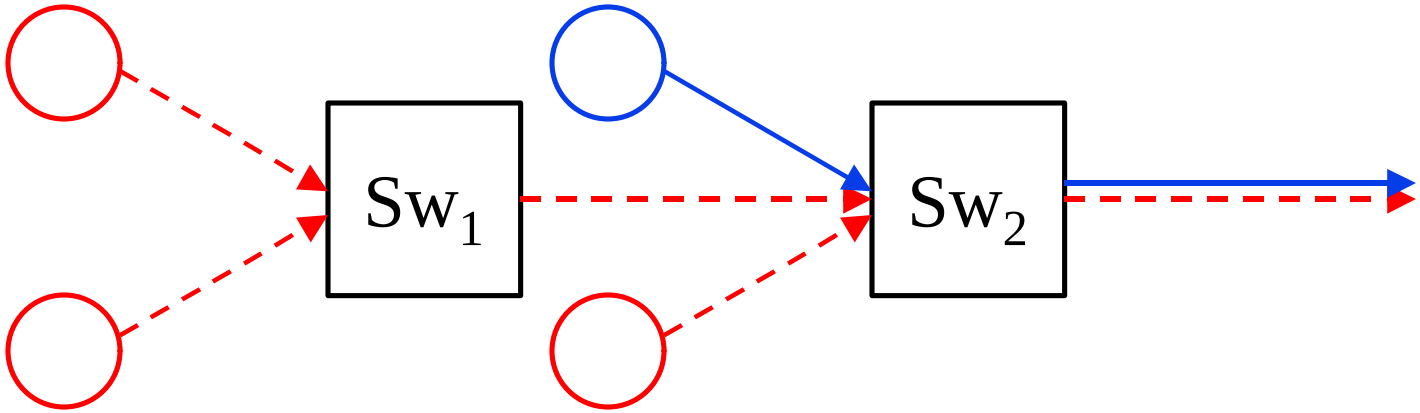

- Sw_2 has to deal with
  - Two VCs: Red and Blue
  - Four sources: Three red sources and one blue source
  - Three flows: Two red flows and one blue

Fairness Definitions

- Source-based: N-to-one connection = N one-to-one connections ⇒ Use max-min fairness among sources
- VC/Source-based:
  1. Allocate bandwidth fairly among VCs
  2. For each VC, allocate fairly among its sources
- Flow-based: Flow = VC coming on an input link. Switch can easily distinguish flows.
- VC/Flow-based:
  1. Allocate bandwidth fairly among VCs
  2. For each VC, allocate fairly among its flows

Mpt-pt Issues and Results

- Cells of senders in the same multipoint-to-point VC cannot be distinguished
- Question: Can we achieve source-based fairness? Answer: Yes!
- We extended ERICA to achieve source based fairness for mpt-pt VCs

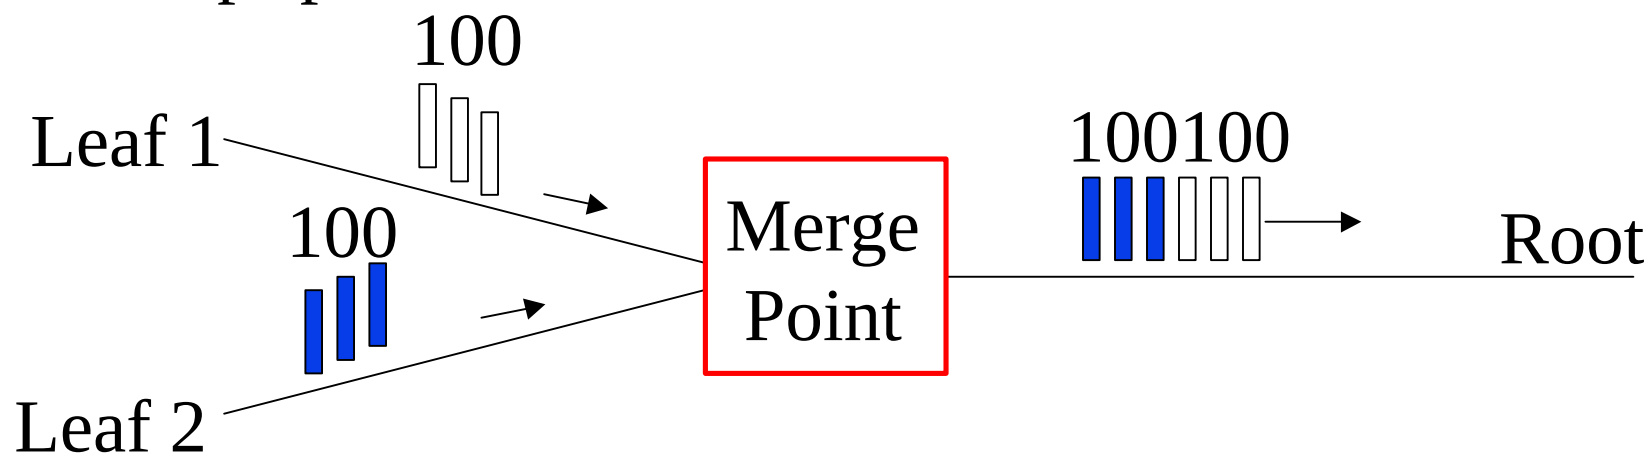

Impact II

- A new item has been added to the living list of TM specs describing the issues in Mpt-to-pt ABR
- A sample merge point algorithm, which applies to mpt-to-mpt also, has been added.
- Our Fairness definitions have been added to TM4.1 spec.

Virtual Source / Virtual Destination (VS / VD)

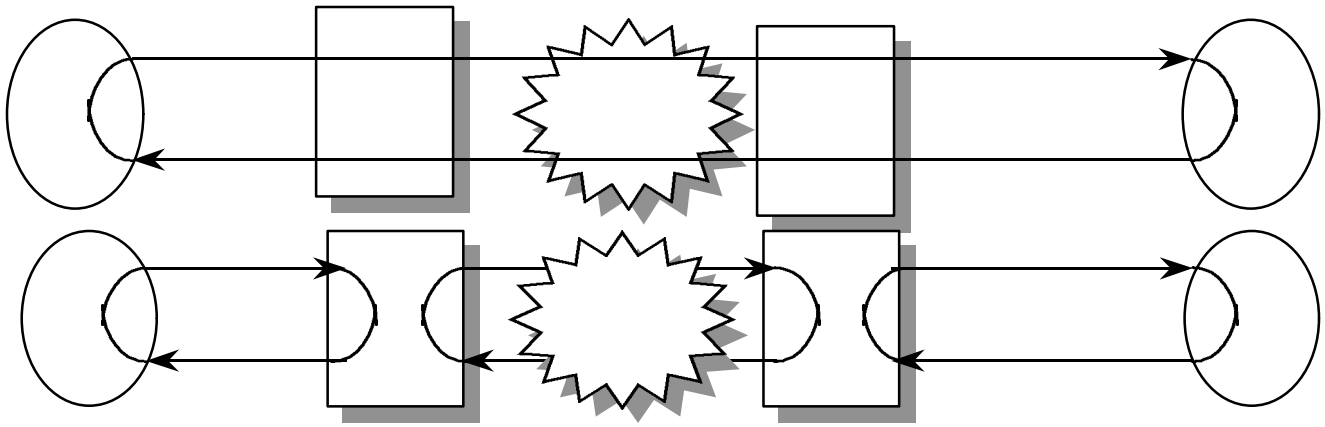

- Segments the end-to-end ABR control loop.
- Coupling between loops is implementation specific.
- ABR switches separated by non-ATM network could also implement VS/VD.

VS/VD Issues

- Although TM4.0 allows VS/VD, it does not describe how the feedback must be passed from VS to previous VD.
- It is not clear if and when VS/VD help.
- Our Accomplishments:
  - Analyzed issues in designing rate allocation schemes for VS/VD switches.
  - Developed a per-VC rate allocation scheme for VS/VD.
  - Showed how VS/VD can help in buffer management across the network.

Simulation Results

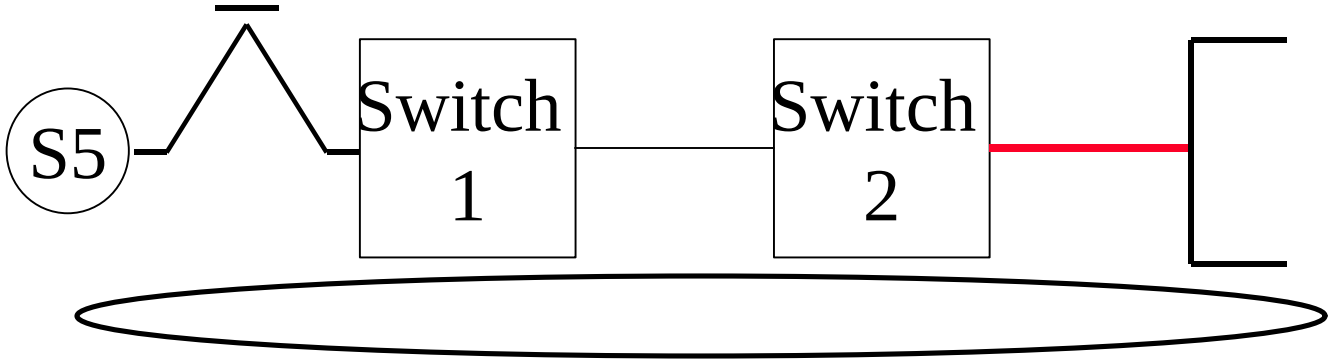

- Without VS/VD:
  - Single control loop for the entire connection.
  - All queues are in the bottleneck switch.
  - Buffer requirements for terrestrial switch are proportional to satellite propagation delay.

Results (Cont)

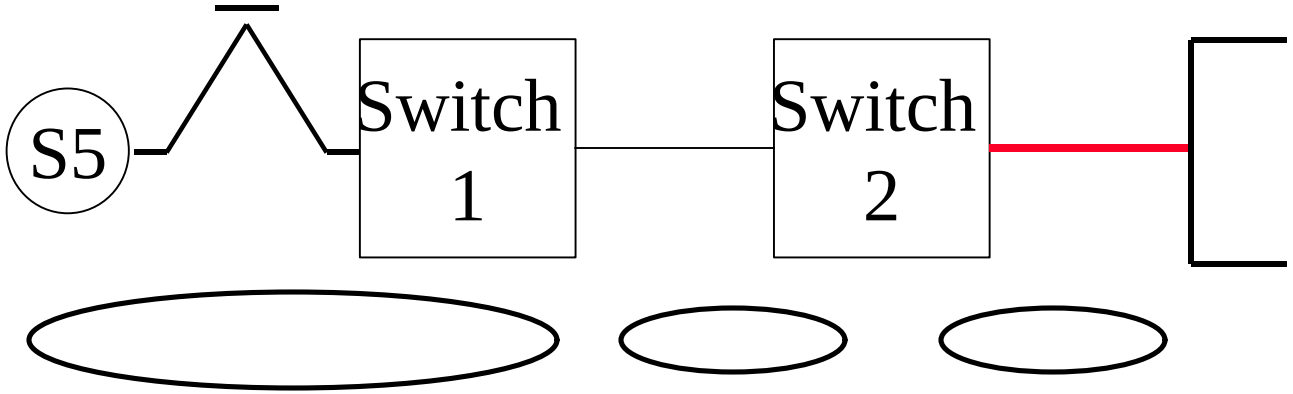

- With correct implementation of VS/VD:
  Maximum queue at each switch
  $\leq$ k* bandwidth delay product of the previous loop
  $\Rightarrow$ Isolate long-delay hops from short-delay hops.
- Workgroup switches on satellite paths will not need buffering proportional to round-trip even if they are the bottleneck.

Impact III

- Our extension of the ERICA switch algorithm including VS/VD was accepted for inclusion as a sample algorithm in the TM4.1 baseline text.

Video over ABR: How?

- Compression parameters can be dynamically adjusted to match the available bandwidth
  $\Rightarrow$ real-time ABR or rt-ABR
- With proper switch algorithm,
  ABR queues in the switches are very small
  $\Rightarrow$ Negligible delay in the network
- Any switch algorithm with fast transient response and queue control can loosely guarantee low delay through the switch

Scheduling and Buffering Issues

- Weighted max-min fairness: Allocate rates to flows in proportion to their weights
  ⇒ Higher rate sources are treated preferentially
- Buffering at the sources and acceptable loss
  ⇒ Equivalent bandwidth
  ⇒ MCR
  ⇒ Minimum acceptable quality is guaranteed
- Internet does not provide MCR. ABR does.
  rt-ABR video will be much better

Summary of Results

- MPEG2 compressed video = piecewise CBR, long-range dependent rate, random inter-MPCR intervals
- ABR with appropriate switch algorithm can handle the randomness in ABR capacity
- With ERICA+ and Infinite TCP Traffic:
  - Queue lengths < 3 × Feedback delay
  - Efficiency close to the maximum possible.
  - Queues are similar to those with deterministic VBR

Generalized Fairness

- Real-time applications need non-zero Minimum Cell Rate (MCR)
- In TM4.0, Distribution of excess bandwidth (fairness) is implementation specific.
- TM4.0 has five examples of fair distribution
- We have shown that two of the examples are not meaningful and have proposed a sixth example that is a superset of the remaining three definitions
- We developed a switch algorithm that implements the proposed definition

Impact IV

The proposed definition was added to TM4.1 baseline text being developed now.

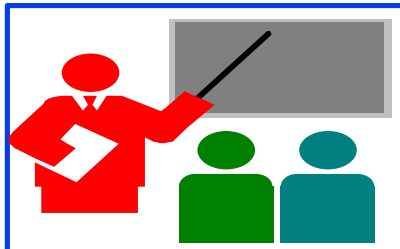

Summary

1. Multipoint Communication: Source-based fairness can be achieved even though sources can not be distinguished in an mpt-pt VC
2. Virtual Source/Virtual Destination: The buffering required at each VS/VD switch is proportional to the bandwidth-delay product of the next loop
3. Real-Time ABR: Generalized Fairness based on charging policies is a superset of TM4.0 policies.
4. Extensions of ERICA to cover the above three cases have been developed

Our Contributions and Papers

- This project resulted in 6 papers and 15 ATM Forum contributions.
- ATM Forum TM4.1 specs (hence, the industry) was impacted as a result of this project.
- All our contributions and papers are available on-line at http://www.cis.ohio-state.edu/~jain/

Thank You!

This work was sponsored in part by
Rome Laboratory/C3BC under
Contract #F30602-96-C-0156