The NS Simulator

CAIDA/IEC Networking Workshop

brought to you by:

Kevin Fall

University of California, Berkeley

NetBoost Corporation

kfall@cs.Berkeley.EDU

http://www.cs.berkeley.edu/~kfall

AND

Polly Huang

ATT Research

formerly USC/ISI

huang@catarina.usc.edu

http://www-scf.usc.edu/~bhuang

Aug 11-12, 1999

Audience and Outline

- Audience
  - educators and network researchers
  - (curious others)
- Topics for today
  - Vint Introduction and Perspective (Kevin)
    * goals, accomplishments, status, architecture
    * state of the software and distributions
  - Essentials/Getting Started (Polly)
  - Objects and the models they represent (Kevin, Polly)
  - Extending the simulator (Polly)
  - Specific Additional Topics (Kevin, Polly)
    * scenario and topology generation
    * wireless LANs, Mobility Support
    * application Level Support
    * emulation
- Topics for tomorrow
  - Examples of NS usage in research and education
    * TCP and SRM studies
  - Exercises: hands-on fun with NS

The Vint Project

- Virtual InterNet Testbed
  - provide a common platform for network research
  - special focus on multi-protocol interactions and scale
  - see `http://netweb.usc.edu/vint`
- Multi-site collaborative simulation environment
  - AT&T Research
  - Lawrence Berkeley National Laboratory
  - UC Berkeley
  - USC/ISI
  - Xerox PARC
- Current contract with DARPA funds project until about Spring 1999

Vint People

- Project leads and co-PIs:
  - Lee Breslau (ATT/PA)
  - Deborah Estrin (USC/ISI)
  - Kevin Fall (UC Berkeley)
  - Sally Floyd (ATT/ACIRI)
  - Mark Handley (ATT/ACIRI)
  - John Heidemann (USC/ISI)
  - Scott Shenker (ATT/ACIRI)
- (Fairly) Recent co-PI moves
  - Kevin Fall from LBNL to UC Berkeley
  - Sally Floyd from LBNL to ATT/ACIRI
  - Scott Shenker from Xerox PARC to ATT/ACIRI
  - Lee Breslau from Xerox PARC to ATT/PA
- Students, graduated students, and ISI staff
  - Nirupama Bulusu
  - Padmaparna Haldar
  - Ahmed Helmy
  - Polly Huang
  - Satish Kumar
  - Yuri Pryadkin
  - Nader Salehi
  - Kannan Varadhan
  - Sam Wu
  - Haobo Yu

Vint Project

- Approach is to support *collaborative* simulation effort
  - promote sharing
  - increase confidence in results
  - establish common reference
- Methods
  - incorporate recent simulation models
  - establish regression test suite
  - current and periodic availability of source code
  - (oh yes, and its free)
- Project Management
  - weekly conferences on the Internet
  - bi-annual (or so) retreats
  - periodic tutorials and workshops
- Simulator base software for Vint is NS2

Development Status

- half-dozen (or so) people actively contributing to the code base
- contributions from: Xerox PARC, USC/ISI, UCB, LBNL, ATT/ACIRI
- Some approximate numbers:
  - 99K lines of C++ code
  - 37K lines of OTcl support code
  - 47K lines of test suites, examples
  - 13K lines of documentation (222 pages in notes/docs)
- Users we know about:
  - 399 entries on ns-users mailing list
  - 274 unique domains represented:
    * 40 COM, 82 EDU, 3 GOV, 5 NET, 4 ORG, 3 MIL
    * 137 non-USA domains, 32 non-USA countries
    * CA, DE, FR, JP, KR, UK each above 12 domains
- See main VINT and NS-2 web pages at:
  `http://netweb.usc.edu/vint`
  `http://www-mash.cs.berkeley.edu/ns/ns.html`
- Open mailing lists: ns-users, ns-announce
  - majordomo@mash.cs.berkeley.edu

Releases

- Releases of the simulator
  - periodic official “releases”
  - nightly “snapshots”
- Releases
  - attempt to do these at important points
  - about every 6 months or so
  - available from Berkeley or UK European mirror
  - current is 2.1b5 (Mar 1999)
- Daily Snapshots
  - snapshot of that day’s source code tree (raw)
  - probably compiles and works, but buyer beware
  - `http://www-mash.cs.berkeley.edu/dist/vint/ns-src-current.tar.gz`
- Platforms
  - most UNIX and UNIX-like things
  - Windows 32
  - (Emulation only for FreeBSD for now)

Building NS

- Two ways to build NS
  - "from the pieces"
  - "all at once"
- From the pieces (you must get these yourself)
  - tcl/tk version 8 (to 8.0.4) (from Scriptics)
  - otcl (version 1.0a4 at UCB)
  - tclcl (version 1.0b8 from UCB; also daily snapshot)
  - ns-2 (source release 2.1b5 from UCB; also daily snapshot)
  - (special ns-2 snapshot with CMU wireless code also available)
  - ns-2 (windows32 binary release)
- Optional pieces
  - nam-1 (version 1.0a7 at UCB; also daily snapshot)
  - xgraph (simple display utility
  - perl (need version 5.003 or later; from perl.com)
  - tcl-debug (vers 1.7 from NIST)
  - dmalloc (vers 3.2.1 from letters.com)
  - topology converters (sgb2ns and tiers2ns) and libraries

Building NS

- All at once
  - pick up `ns-allinone-2.1b5a.tar.gz` from UCB
  - run the install script
- In any case, check the bugs web page:
  - `http://www-mash.cs.berkeley.edu/ns/ns-problems.html`

NS Validation

- Validation of simulation models
  - much of the functionality has a *validation test*
  - the `validate` script runs all validation tests
  - validation of specific functionality can also be performed
  - (see the tcl/test directory for the tests)
- A word of caution (from the main web page):

  While we have considerable confidence in ns, ns is not a polished and finished product, but the result of an on-going effort of research and development. In particular, bugs in the software are still being discovered and corrected. Users of ns are responsible for verifying for themselves that their simulations are not invalidated by bugs. We are working to help the user with this by significantly expanding and automating the validation tests and demos. Similarly, users are responsible for verifying for themselves that their simulations are not invalidated because the model implemented in the simulator is not the model that they were expecting. The ongoing ns Notes and Documentation should help in this process.

Contributions

- Contributed code
  - we accept code contributions into the simulator
  - we also will point to your work on our contrib page
- Requirements for the core
  - source code
  - manual/documentation description
  - validation test suites
- Other code
  - contributed code web page
  - point us to your work
  - we are considering removing code lacking validation tests
- Research results
  - we would like to know about research you've done using NS
  - we will add a link to your work at the research web page

NS2 Architecture

- Object-oriented structure
  - evolution from NSv1 (split tcl/C++ objects)
  - OTcl: object-oriented extension to Tcl
    (from David Wetherall at MIT/LCS for VuSystem)
    (now supported by UCB Mash group)
- Control/“Data” separation
  - control operations in OTcl (for easy configuration)
  - data pass through C++ objects (for speed)
- Modular approach
  - fine-grain object decomposition
  - positives: composible, re-usable
  - negatives: must “plumb” in OTcl,
    developer must be comfortable with both environments,
    tools (fairly steep learning curve)

And now Polly...

The Simulator

- Simulator API is a set of methods belonging to a *simulator* object:
- Create a simulator with:

```
set ns [new Simulator]
```

- What this does:
  - initialize the packet format (calls `create_packetformat`)
  - create a scheduler (defaults to using a calendar queue)
  - select the default address format (later)
- Scheduler:
  - handles time, timers and events (packets), deferred executions ("ATs")
  - `Scheduler/List` - linked-list scheduler
  - `Scheduler/Heap` - heap-based scheduler
  - `Scheduler/Calendar` - calendar-queue scheduler
  - `Scheduler/RealTime` - real-time (for emulation)
  - see Reeves, "Complexity Analyses of Event Set Algorithms", *The Computer Journal*, 27(1), 1984

Using the scheduler

- Scheduler API is through Simulator object:

  `Simulator instproc now` ;# *return scheduler's notion of current time*
  `Simulator instproc at args` ;# *schedule execution of code at specified time*
  `Simulator instproc at-now args` ;# *schedule execution of code now*
  `Simulator instproc after n args` ;# *schedule exec of code after n secs*
  `Simulator instproc run args` ;# *start scheduler*
  `Simulator instproc halt` ;# *stop (pause) the scheduler*
  `Simulator instproc create-trace type files src dst` ;# *create trace object*
  `Simulator instproc create_packetformat` ;# *set up the simulator's packet format*

- Example:

```
MySim instproc begin {} {
        ...
        set ns_ [new Simulator]
        $ns_ use-scheduler Heap
        $ns_ at 300.5 "$self complete_sim"
        ...
}
MySim instproc complete_sim {} {
        ...
}
```

Simulator Timing

- each topology object has a generic receive method
  `NsObject::recv(Packet*, Handler* h = 0)`
- most objects have single neighbor `Connector::target_`
- *cut-through* transfers: send packet directly to neighbor without involving scheduler
  `Connector::send(Packet* p) { target_->recv(p); }`
- *scheduling barriers*:
  - any point that advances time into future
    (i.e., delay element)
  - need inter-object "protocol" to decouple timing
  - barrier takes non-null Handler
  - schedule delay and invoke handler on completion
  - example: queue/delay objects (later)

Mathematical Support

- Random number generation
  - RNG implemented in simulator
    (should produce same results on various platforms)
  - based on S. Park and K Miller, CACM 31:10, Oct. 1988
  - support for multiple streams
  - different seeding options
- Random variables
  - distributions applied to RNG streams
  - distributions: uniform, exponential, pareto, constant, hyper-exponential, normal, log-normal
  - empirical (driven from trace file, list of values)
- Integrals
  - approximation of integral by discrete sums
  - used for average queue size computations
- Samples
  - collect samples
  - provides mean, variance, sum, and count

Applications

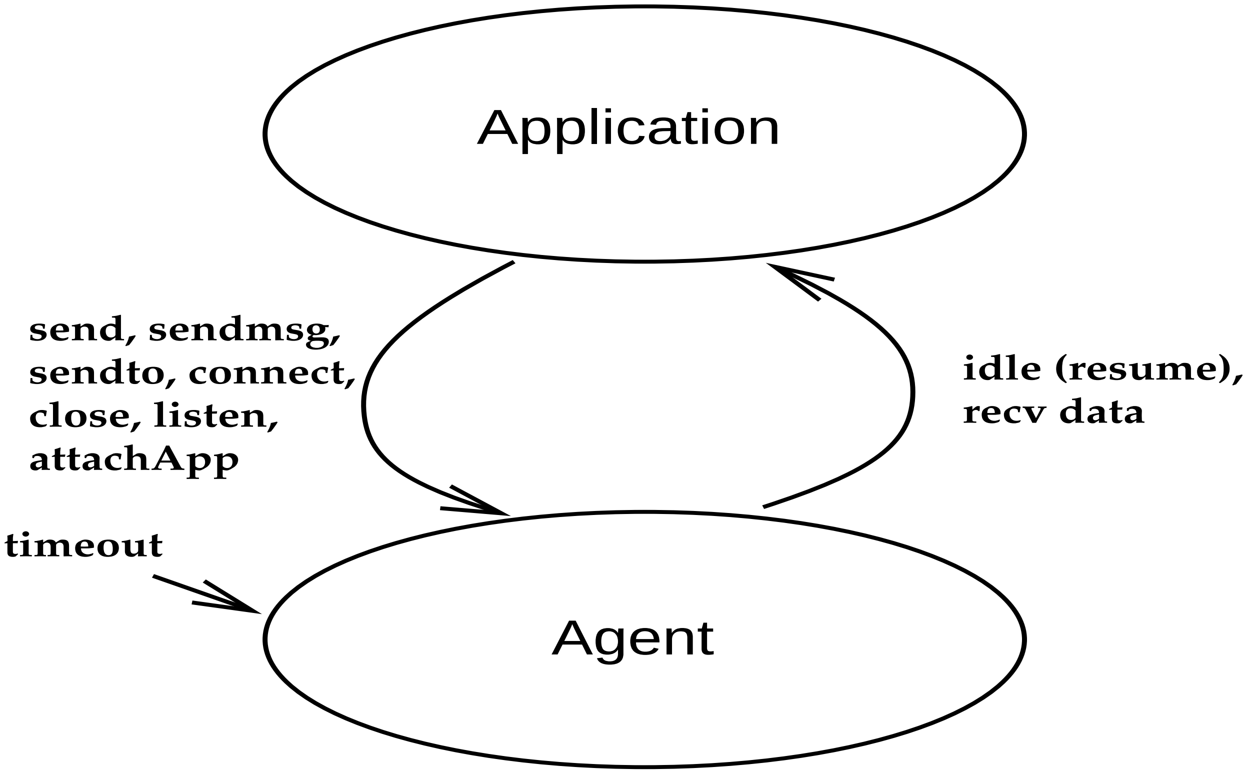

- *Application:* a model of an application, usually a traffic source
- has an associated *agent*, which usually corresponds to a transport entity
- interface is somewhat like *sockets*

Sources and Traffic Generation

- Applications are of two types: *sources* and *traffic generators*
- Sources are used to drive stream transports (e.g. TCP)
- Traffic generators are used to drive connectionless transports (e.g. UDP)
- Sources
  - Telnet - simulates characters typed by a user
  - FTP - bulk data transfer
- Traffic Generators
  - EXPOO - exponential on/off times, sent at fixed rate
  - POO - pareto on/off times, sent at fixed rate
  - CBR - deterministic rate
  - TrafficTrace - use trace file containing time/len pairs

Sources

- class `Application/Telnet`
  - may specify `interval_` parameter
  - if zero, picks randomly among 10000 measured interarrivals (TCPLIB)
  - if nonzero, uses scaled exponential for interarrivals
  - packet size constant (but available via bind call)
  - Methods:
    - **`start`** - continuous send
    - **`stop`** - stop sending
    - **`send`** n - send n bytes
- class `Application/FTP`
  - bulk data sender used to drive TCP
  - implemented entirely in OTcl
  - Methods:
    - **`start`** - continuous send
    - **`stop`** - stop sending
    - **`send`** n - send n bytes
    - **`produce`** n - send n packets
    - **`producemore`** n - send n more packets

Traffic Generation

- generate traffic according to distributions or traces
- generally used to drive CBR or UDP transport agents
- Exponential (class `Application/Traffic/Exponential`)
  - exponentially distributed on/off times
  - parameters: ontime, offtime, rate, packet size
  - what these mean:
    * burst for expo time with mean ontime
    * be silent for expo time with mean offtime
    * while bursting, send at rate rate
    * use appropriate inter-departure time given rate/size
- Pareto (class `Application/Traffic/Pareto`)
  - pareto distributed on/off times
  - (many aggregated together can be LRD)
  - parameters: ontime, offtime, rate, shape, packet size
  - what these mean:
    * like expo, except pareto using shape parameter

Trace-Based Traffic Generator

- generate traffic according to trace file
- two classes: `Tracefile` and `Application/Traffic/Trace`
- trace file uses small binary format:
  - first 32-bit field: inter-packet time (microsecs)
  - second 32-bit field: packet size (bytes)

The NS Address Space

- Addresses are globally unique, and comprise two parts:
  - node id
  - port id
- n bits are available to hold both:
  - $n \leq 32$, with default 16
  - divided into (node ID, port ID)
  - node ID globally unique, port ID locally significant
  - (so, it mixes network and transport layers)
- limitations of the default
  - upper bit of node ID is for multicast addresses
  - thus, $\rightarrow$ **128 max nodes**, **256 max agents on a node**
- default can be a problem, so can *expand* the space

Expanding the Address Space

- The *expanded* address space does this:
  - 30 bits to hold (node ID, port ID)
  - 1 multicast bit, 21 bits for node ID
  - 8 bits port ID
- Turn it on with:
  - `$ns set-address-format expanded`
- You can also expand the port space (up to some max $\leq n$)
  - `$ns expand-port-field-bits` *nbits*
  - expands the number of bits allocated for port numbers
  - allows for more than 256 agents on a node

Hierarchical Addressing

- Addresses may be placed into a hierarchy
  - default configuration
  - specified configuration
  - Other facts:
    * enabling this enables hierarchical routing
    * always leaves 8 bits for port ID
- Default hierarchy
  - 24 bits of node ID in 3 levels of hierarchy
  - 8 bits per level of hierarchy
  - 8 bits of port ID
  - with multicast enabled, one less bit at top level
  - `$ns set-address-format hierarchical`
- Specified hierarchy
  - specify explicit number of levels
  - also specify number of bits for each level
  - `$ns set-address-format hierarchical 3 6 8 8`
  - creates 3 levels of hierarchy

Agents

- *Agents:* usually a transport protocol endpoint/entity (but may also be used for implementing routing protocols)
- Where they fit in:

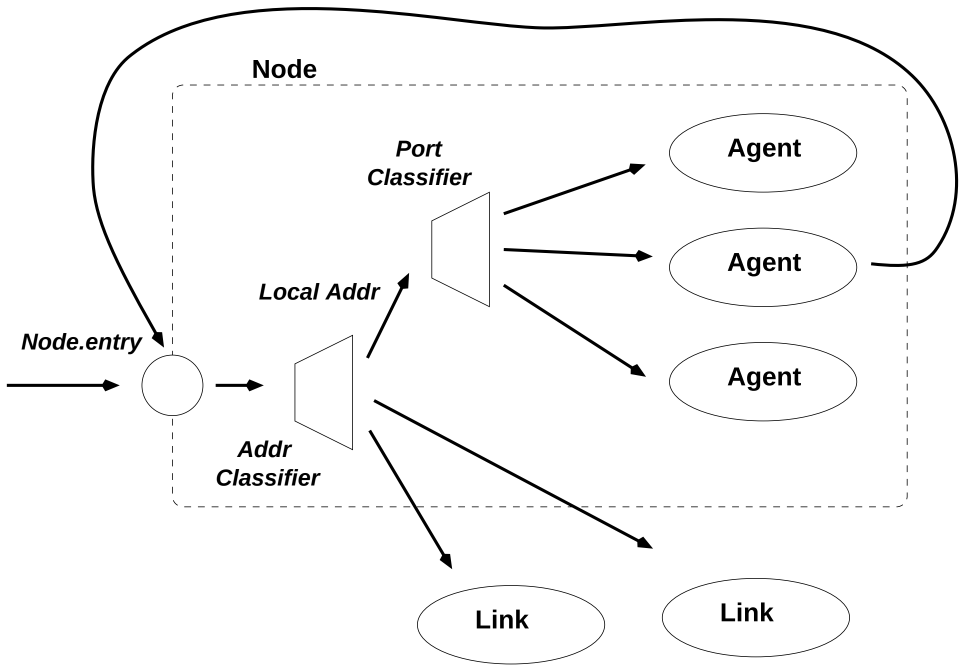

- What they provide:
  - a local and destination port address
  - functions for helping to generate/fill-in in packet fields
  - a base class supporting the application interface

Creating a new Agent

- The `Agent` class:

```
class Agent : public Connector {
 public:
        Agent(int pktType);
        virtual ~Agent();
        virtual void timeout(int tno);
        virtual void sendmsg(int nbytes, const char *flags = 0);
        virtual void send(int nbytes) { sendmsg(nbytes); }
        virtual void sendto(int nbytes, const char* flags, nsaddr_t dst);
        virtual void connect(nsaddr_t dst);
        virtual void close();
        virtual void listen();
        virtual void attachApp(Application* app);
        virtual int& size() { return size_; }
        inline nsaddr_t& addr() { return addr_; }

 protected:
        int command(int argc, const char*const* argv);
        void recv(Packet*, Handler*);
        ...
```

- basic tasks to create a new agent:
  1. decide its inheritance structure
  2. create the class and fill in the API virtual functions
  3. define OTcl linkage functions (later)
  4. write the necessary OTcl code to access your agent
- hardest part may be understanding the OTcl/C++ interaction (fortunately, much of this is shielded from you if you so choose)

Example: the Message Agent

- provides a very simple place to store a message
- Packet header (from **message.h**):

```
struct hdr_msg {
        char msg_[64];
        static int offset_;
        inline static int& offset() { return offset_; }
        inline static hdr_msg* access(Packet* p) {
                return (hdr_msg*) p->access(offset_);
        }
        /* per-field member functions */
        char* msg() { return (msg_); }
        int maxmsg() { return (sizeof(msg_)); }
};
```

- OTcl linkage (for class creation, from **message.cc**):

```
int hdr_msg::offset_;
static class MessageHeaderClass : public PacketHeaderClass {
public:
        MessageHeaderClass() :
           PacketHeaderClass("PacketHeader/Message",
                                   sizeof(hdr_msg)) {
               bind_offset(&hdr_msg::ofset_);
        }
} class_msghdr;
```

Example: the Message Agent (cont)

- The class definition, constructor and variable linkage:

```
class MessageAgent : public Agent {
 public:
        MessageAgent();
        int command(int argc, const char*const* argv);
        void recv(Packet*, Handler*);
};
static class MessageClass : public TclClass {
public:
        MessageClass() : TclClass("Agent/Message") {}
        TclObject* create(int, const char*const*) {
                return (new MessageAgent());
        }
} class_message;

MessageAgent::MessageAgent() : Agent(PT_MESSAGE)
{
        bind("packetSize_", &size_);
}
```

Example: the Message Agent (cont)

- Main functions:

```
void MessageAgent::recv(Packet* pkt, Handler*)
{
     hdr_msg* mh = hdr_msg::access(pkt);
     ... process packet ...
     Packet::free(pkt);
}
int MessageAgent::command(int argc, const char*const* argv)
{
     Tcl& tcl = Tcl::instance(); // call into interp
     if (argc == 3) { // $obj send msgtext
          if (strcmp(argv[1], "send") == 0) {
               Packet* pkt = allocpkt();
               hdr_msg* mh = hdr_msg::access(pkt);
               const char* s = argv[2];
               int n = strlen(s);
               if (n >= mh->maxmsg()) {
                       tcl.result("message too big");
                       Packet::free(pkt);
                       return (TCL_ERROR);
               }
               strcpy(mh->msg(), s);
               send(pkt, 0);
               return (TCL_OK);
          }
     }
     return (Agent::command(argc, argv)); // for inheritance
}
```

UDP and RTP

- UDP Agent:
  - non-connection-oriented sending agent
  - fixed size packets
  - has some RTP-like additions:
    * numbers packets sequentially
    * marks send time in packets
    * marks burst ("talkspurt") endpoints
- RTP Agent:
  - similar function as UDP
  - also includes linkage to RTPSession Object
- RTP Session and RTP Source Objects:
  - provides functions like RTCP (RFC1889)
  - RTPSource object provides stats about an RTP sender
  - RTPSession demuxes received reports and generates them (as needed by RTCP)

TCP Agents

- Two categories: one-way and two-way ("full TCP")
- One-way variants of TCP available:
  - Agent/TCP - a "tahoe" TCP sender
  - Agent/TCP/Reno - a "Reno" TCP sender
  - Agent/TCP/NewReno - Reno with a modification
  - Agent/TCP/Sack1 - TCP with selective repeat (follows RFC2018)
  - Agent/TCP/Vegas - TCP Vegas
  - Agent/TCP/Fack - Reno TCP with "forward acknowledgement"
  - Agent/TCP/Session - shared congestion state w/multiple connections
  - Agent/TCP/Int - per-connection reliability for use w/Session
  - Agent/TCP/*/RBP - Reno, Vegas with Rate Based Pacing
  - Agent/TCP/Asym - TCP mods for asymmetric channels
- One-way TCP receiving agents currently supported are:
  - Agent/TCPSink - one ACK per packet
  - Agent/TCPSink/DelAck - configurable delay per ACK
  - Agent/TCPSink/Sack1 - selective ACK sink (follows RFC2018)
  - Agent/TCPSink/Sack1/DelAck - Sack1 with DelAck
  - Agent/TCPSink/Asym - sink for Asym senders

Two-Way TCP Agents

- Two-way TCP agents:
  - Agent/TCP/FullTcp - provides Reno functionality
  - Agent/TCP/FullTcp/Tahoe (newish)
  - Agent/TCP/FullTcp/Sack (newish)
- One-way and two-way TCPs are not interoperable

Base One-Way TCP Agents

- `TCP` (Tahoe), `TCP/Reno`, and `TCP/NewReno`
- Common features:
  - computations all in packet units w/configurable packet size
  - fast retransmit
  - slow-start and congestion avoidance
  - dynamic RTT estimation and RTX timeout assignment
  - simulated (constant) receiver's advertised window
- Tahoe TCP:
  - perform slow-start on any loss (RTO or fast retransmit)
  - no fast recovery
- Reno TCP:
  - fast *recovery*: inflate *cwnd* by dup ack count until new ACK
  - slow-start on RTO
  - on fast retransmit:

    $$cwnd \leftarrow curwin/2 \ , \ ssthresh \leftarrow cwnd$$
- "Newreno" TCP:
  - modest modification to Reno TCP
  - only exit fast recovery after ACK for highest segment arrives
  - helps reduce "stalling" due to multiple packet drops in a window

Other TCP Agents

- `TCP/Sack`, `TCP/Fack`, and `TCP/Vegas`
- Selective ACK TCP:
  - SACK simulation based on RFC2018
  - ACKs carry extra information indicating received segments
  - requires SACK-aware sink
  - sender avoids sending redundant info
  - default to 3 "SACK blocks" (for using timestamps, see RFC2018)
    - block contains start/end sequence numbers
    - block containing most recently received segment always present
  - regular ACK number still gives final say
- Fack TCP:
  - "forward ACK" TCP (experimental, see SIGCOMM '96)
  - use SACK info for estimate of packets in the network
  - overdamping algorithm (to limit slow-start overshoot)
  - rampdown algorithm (for transmission smoothing)
- Vegas TCP:
  - contributed code from Ted Kuo (NC State Univ)

TCP Agent Parameters

- Common configuration parameters and defaults for TCP agents:

```
Agent/TCP set maxburst_ 0                  ;# max pkts emitted due to 1 recvd
Agent/TCP set maxcwnd_ 0                   ;# max bound on congestion window
Agent/TCP set syn_ false                   ;# do SYN exchange prior to data xfer
Agent/TCP set tcpip_base_hdr_size_ 40      ;# size of TCP/IP hdr, no opts
Agent/TCP set timestamps_ false            ;# due RFC1323-style time stamps
Agent/TCP set window_   20                 ;# max bound on window size
Agent/TCP set windowInit_ 1                ;# initial/reset value of cwnd
Agent/TCP set windowOption_ 1              ;# cong avoid algorithm (1: standard)
Agent/TCP set windowConstant_ 4            ;# used only when windowOption != 1
Agent/TCP set windowThresh_ 0.002          ;# used in computing averaged window
Agent/TCP set overhead_ 0                  ;# !=0 adds random time between sends
Agent/TCP set ecn_ 0                       ;# TCP should react to ecn bit
Agent/TCP set packetSize_ 1000             ;# packet size used by sender (bytes)
Agent/TCP set bugFix_ true                 ;# see documentation
Agent/TCP set slow_start_restart_ true     ;# do slow-start after idle period
Agent/TCP set tcpTick_ 0.1      ;# timer granularity in sec (.1 is NONSTANDARD)
Agent/TCP set maxrto_ 100000               ;# bound on RTO (seconds)
Agent/TCP set srtt_init_ 0                 ;# initial value for smoothed rtt est
Agent/TCP set rttvar_init_ 12              ;# initial value for rtt var est
Agent/TCP set rtxcur_init 6.0              ;# initial value for current rtx timer
Agent/TCP set T_SRTT_BITS 3                ;# # bits after binary point for SRTT
Agent/TCP set T_RTTVAR_BITS 2              ;# # bits after binary point for RTTVAR
Agent/TCP set rttvar_exp_ 2                ;# exponent of 2 which multiples rttvar
```

- Dynamic values of interest:

```
Agent/TCP set dupacks_ 0                   ;# duplicate ACK counter
Agent/TCP set ack_ 0                       ;# highest ACK received
Agent/TCP set cwnd_ 0                      ;# congestion window (packets)
Agent/TCP set awnd_ 0                      ;# averaged cwnd (experimental)
Agent/TCP set ssthresh_ 0                  ;# slow-stat threshold (packets)
Agent/TCP set rtt_ 0                       ;# rtt sample
Agent/TCP set srtt_ 0                      ;# smoothed (averaged) rtt
Agent/TCP set rttvar_ 0                    ;# mean deviation of rtt samples
Agent/TCP set backoff_ 0                   ;# current RTO backoff factor
Agent/TCP set maxseq_ 0                    ;# max (packet) seq number sent
```

TCP Sink Agents

- Sinks for one-way TCP senders
- Types
  - standard sinks, delayed-ACK sinks, SACK sinks
- Standard sinks:
  - generate one ACK per packet received
  - ACK number overloaded in "sequence number" packet field
- Delayed-ACK sinks:
  - same as standard, but with variable delay added between ACKs
  - time to delay ACKs specified in seconds
- SACK sinks:
  - generates additional information for SACK capable sender
  - configurable `maxSackBlocks_` parameter
- Asym sinks:
  - ACK pacing

Two-Way TCP ("FullTCP")

- most TCP objects are one-way (and require a source/sink pair)
- real TCP can be bi-directional
- simultaneous two-way data transfer alters TCP dynamics considerably
- the `TCP/FullTcp` agent:
  - follows closely to "Reno" TCP implementation in 4.4 BSD
  - byte-oriented transfers
  - two-way data supported
  - most of the connection establishment/teardown
  - symmetric: only one agent type used for both sides
- Differences from the "real thing":
  - no receiver's advertised window/persist mode
  - no urgent pointer
  - no 2MSL wait
  - no RST segments
- Now supports Tahoe, NewReno, and Sack variants

FullTCP Parameters

- Parameters and defaults:

```
Agent/TCP/FullTcp set segsperack_ 1      ;# segs received before generating ACK
Agent/TCP/FullTcp set segsize_ 536      ;# segment size (MSS size for bulk xfers)
Agent/TCP/FullTcp set tcprexmtthresh_ 3   ;# dupACKs thresh to trigger fast rexmt
Agent/TCP/FullTcp set iss_ 0            ;# initial send sequence number
Agent/TCP/FullTcp set nodelay_ false     ;# disable sender-side Nagle algorithm
Agent/TCP/FullTcp set data_on_syn_ false        ;# send data on initial SYN?
Agent/TCP/FullTcp set dupseg_fix_ true    ;# avoid fast rxt due to dup segs+acks
Agent/TCP/FullTcp set dupack_reset_ false  ;# reset dupACK ctr on !0 len data seg
s containing dup ACKs
Agent/TCP/FullTcp set interval_ 0.1            ;# delayed ACK interval
Agent/TCP/FullTcp set close_on_empty_ false    ;# close conn after send all data
Agent/TCP/FullTcp set ts_option_size_ 10     ;# size of rfc1323 ts option (bytes)
Agent/TCP/FullTcp set reno_fastrecov_ true     ;# congestion window inflation
Agent/TCP/FullTcp set pipectrl_ false      ;# use "pipe" model congestion ctrl
Agent/TCP/FullTcp set open_cwnd_on_pack_ true   ;# increase cwnd on partial acks

Agent/TCP/FullTcp/Newreno set recov_maxburst_ 2   ;# maxburst during recovery

Agent/TCP/FullTcp/Sack set sack_block_size_ 8    ;# # bytes in a SACK block
Agent/TCP/FullTcp/Sack set sack_option_size_ 2        ;# # bytes in opt hdr
Agent/TCP/FullTcp/Sack set max_sack_blocks_ 3      ;# # max # of sack blks
```

Other Simple Agents

- the `LossMonitor` agent:
  - monitors arrivals of packets
  - looks for sequence number holes
  - provides counters for:
    - **`nlost_`** - number of holes in number space
    - **`npkts_`** - packet arrivals
    - **`bytes_`** - byte arrivals
    - **`lastPktTime_`** - time of last arrival
    - **`expected_`** - next seq number expected
- the `Message` agent:
  - very simple agent
  - allows for including text "messages" in packets
  - currently limited to at most 64 byte (short) messages

Connectors

- *Connector:* simple in/out topology object with "drop target"

```
/*
 * An NsObject with only a single neighbor.
 */
class Connector : public NsObject {
public:
        Connector();
        inline NsObject* target() { return target_; }
        virtual void drop(Packet* p);
protected:
        int command(int argc, const char*const* argv);
        void recv(Packet*, Handler* callback = 0);
        inline void send(Packet* p, Handler* h) {
                target_->recv(p, h);
        }

        NsObject* target_;
        NsObject* drop_;  // dest for drops
};
```

- if drop target undefined, dropped packets are freed

Introducing Errors

- Packet errors may be introduced into the topology
  - *Error Module* - collection of error models
  - *Error Model* - determines types of errors
- `Error Module` capabilities:
  - insert *classifier* for per-flow error control
  - insert multiple error models
  - collects dropped packets from each error model
  - *default* entry indicates where non-matching traffic goes
- Example:

```
set lossylink_ [$ns link $node_(r1) $node_(k1)]
set em [new ErrorModule Fid]
set errmodel [new ErrorModel/Periodic]
$errmodel unit pkt
$errmodel set offset_ 1.0
$errmodel set period_ 25.0
$lossylink_ errormodule $em
$em insert $errmodel
$em bind $errmodel 0
```

Error Models

- *Error Model:* a parameterized lossy connector
  (can be used as a base class for other loss models)
- drops packet, sets "error" bit, or ECN indication
- error *units*: packets, bits, time
- base version is associated with random variable or periodic parameters:

```
set errmodel [new ErrorModel]
$errmodel unit pkt
$errmodel set rate_ 0.01
$errmodel ranvar [new RandomVariable/Uniform]
```

- if drop target undefined, sent to module drop target
- Specialized Models:
  - `TwoState` – error and error-free Markov model
  - `MultiState` – arbitrary state transitions in Tcl
  - `Periodic` – sinusoidal (period, phase, burst length)
  - `List` – specifies individual packets/bytes
  - `Select` – like Periodic, but uses packet uid's
  - `SRM` – like Select, but for SRM data only
  - `Trace` – like List, but uses a file
  - `Mroute` – affects certain types of mcast prune messages

Queue Management and Packet Scheduling

- *buffer management:* how to hold and toss (mark) packets
- *packet scheduling:* what packets get to depart when
- Buffer management:
  - Drop-tail (FIFO)
  - Random Early Detection (RED)
- Packet scheduling:
  - FIFO
  - CBQ (includes priority + round-robin)
  - Round-robin (DRR)
  - Variants of FQ (FQ, SFQ, WFQ not in base distribution)

Queue Handlers

- Dequeued packets are often sent downstream to *delays*
- delays (usually) cause two actions:
  1. the packet is scheduled to arrive downstream at time $t + d$
  2. the queue becomes unblocked at time t
  3. t is transmit time, d is prop delay time
- so, delays represent a commonly-occurring *scheduling barrier*
- Queue parameters:

```
Queue set limit_ 50                ;# max packet count in queue
Queue set blocked_ false           ;# queue starts off blocked
Queue set unblock_on_resume_ true      ;# queue is unblocked after resume
```

- control of blocking can be useful for queue banks (e.g. CBQ)

Drop Tail and RED Queues

- Drop-Tail Queues (`Queue/DropTail` class)
  - simple FIFO, drop-tail queues
  - drop from tail when occupancy reaches `qlim_`
- RED (Random Early Detection) Queues (`Queue/RED` class)
  - *active* buffer management technique
  - two thresholds: *minth* and *maxth*
  - also a maximum probability *maxprob*
  - compute *average* queue occupancy over time
  - if average exceeds maxth (or `qlim_`) drop a packet
  - if average is under minth, allow packet to enter queue
  - between, scale drop probability linearly on $[0, maxprob]$

RED Queue Parameters

- `bytes_` do computations in bytes instead of packets (requires assignment of a mean packet size estimate)
- `thresh_` - min thresh
- `maxthresh_` - max thresh
- `mean_pktsize_` - used for computing estimated link utilizations during idle periods
- `q_weight_` - weight given to instantaneous queue occupancy for EWMA
- `wait_` - RED should force a wait between drops
- `linterm_` - reciprocal of maxprob
- `setbit_` - mark instead of drop
- `drop-tail_` - drop new pkt instead of random one

CBQ: Class Based Queuing

- Floyd and Jacobson, "Link-sharing and Resource Management Models for Packe t Networks", ToN, Aug 1995
- rewrite from CBQ code in ns-1
- packets are members of *classes*
- classes may contain a *priority* and a *bandwidth allocation*
- classes may *borrow* unused bandwidth from other classes
- packets are scheduled using a round-robin scheduler according t the classes they belong to:
  - packet-by-packet RR
  - weighted RR
  - high-to-low priority

CBQ Implementation

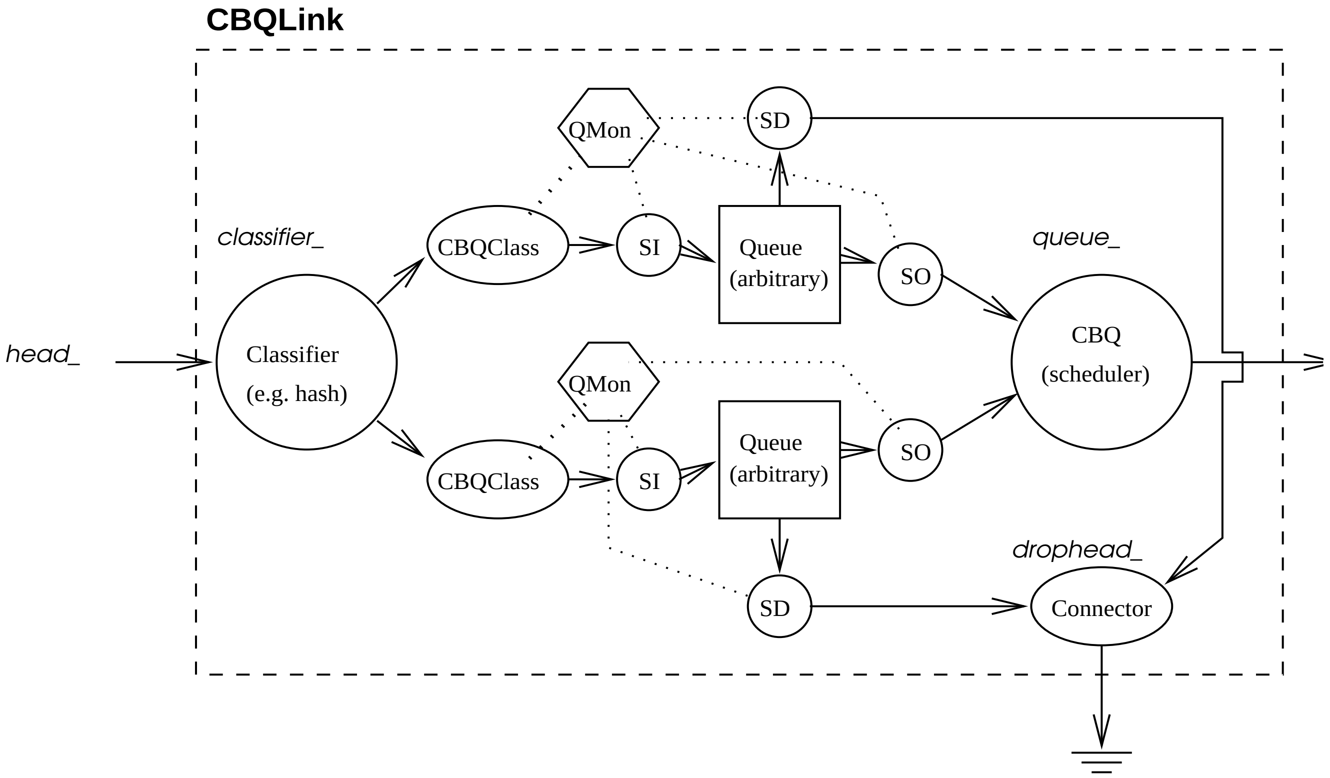

- Major components:
  - classifier (maps packets to classes)
  - classes (holds class state)
  - scheduler (schedules packet departures)
- Implemented as a subclass of link: *CBQ link*

Token Bucket Filters

- Traffic regulation/conditioning by *token bucket filters*
- `TBF` class provides this:
  - `rate_` - token generation rate (ρ)
  - `bucket_` - bucket depth (σ)
  - `qlen_` - maximum queue length (for use as policer)
- Functions:
  - `$agent attach-tbf` *tbf* - attach tbf after agent
  - `$sim attach-tbf-agent` *node agent tbf* - attaches agend to node, then calls attach-tbf

Trace and Monitoring Support

- Two main items: *traces* and *monitors*
- Traces - write an entry for some event (often packet arrivals/departures/drops)
  - `Trace/Enque` - a packet arrival (usually at a queue)
  - `Trace/Deque` - a packet departure (usually at a queue)
  - `Trace/Drop` - packet drop (packet delivered to drop-target)
- Monitors - keep statistics about arrivals/departures/drops (and flows)
  - `SnoopQueue/Out` - on output, collect a time/size sample (pass pac ket on)
  - `SnoopQueue/Drop` - on drop, collect a time/size sample (pass pack et on)
  - `SnoopQueue/EDrop` - on an "early" drop, collect a time/size sampl e (pass packet on)
  - `QueueMonitor` - receive and aggregate collected samples from snoo pers
  - `QueueMonitor/ED` - queue-monitor capable of distinguishing between "early" and standard packet drops
  - `QueueMonitor/ED/Flowmon` - per-flow statistics monitor (manager)
  - `QueueMonitor/ED/Flow` - per-flow statistics container

Trace File Format

- File format for traces generally of this form:

```
+ 1.45176 2 3 tcp 1000 ---- 1 256 769 27 48
+ 1.45276 2 3 tcp 1000 ---- 1 256 769 28 49
- 1.46176 2 3 tcp 1000 ---- 1 256 769 22 43
+ 1.46176 2 3 tcp 1000 ---- 1 256 769 29 50
+ 1.46276 2 3 tcp 1000 ---- 1 256 769 30 51
d 1.46276 2 3 tcp 1000 ---- 1 256 769 30 51
- 1.47176 2 3 tcp 1000 ---- 1 256 769 23 44
+ 1.47176 2 3 tcp 1000 ---- 0 0 768 3 52
+ 1.47276 2 3 tcp 1000 ---- 0 0 768 4 53
d 1.47276 2 3 tcp 1000 ---- 0 0 768 4 53
```

- Fields: arrival/departure/drop, time, trace link endpoints, packet type, size, flags, flow ID, src addr, dst addr, sequence number, uid
- Many of these fields are from the common packet header:

```
struct hdr_cmn {
        int     ptype_;         // packet type (see above)
        int     size_;          // simulated packet size
        int     uid_;           // unique id
        int     error_;         // error flag
        double  ts_;            // timestamp: for q-delay measurement
        int     iface_;         // receiving interface (label)
        int     direction_;     // direction: 0=none, 1=up, -1=down
        int     ref_count_;     // free the pkt until count to 0
        ...
};
```

Trace File Flags

- 7th field of trace file is a set of flags
- see the files `flags.h` and `trace.cc`
- in order they are:
  1. `C` - ECN Echo
  2. `P` - Priority marking (not used)
  3. `-` - (not used)
  4. `A` - Congestion action (endpoint reaction)
  5. `E` - Congestion experienced along path (set by RED)
  6. `F` - Fast start mark (see tcp-fs)
  7. `N` - ECN Capable Transport sent packet

Trace Callbacks

- may opt to invoke a Tcl function in lieu of writing to file
- see the file `tcl/ex/callback_demo.tcl`

```
MyTest instproc begin {} {
        ...
        $link12_ trace-callback $ns_ "$self dofunc"
        ...
}

MyTest instproc dofunc args {
        ... process args ...
}
```

- Args passed to the callback are a string containing a trace output line (e.g.):
  `- 0.80612 0 1 tcp 1000 ------ 0 0.0 1.0 9 13`

Monitors

- Queue monitors: aggregation points for arrival/depart/drop stats
- Flow monitors: similar, but on a per-flow basis
- Snoop queues: part of the topology, "taps" packet flow, delivers samples to associated monitor

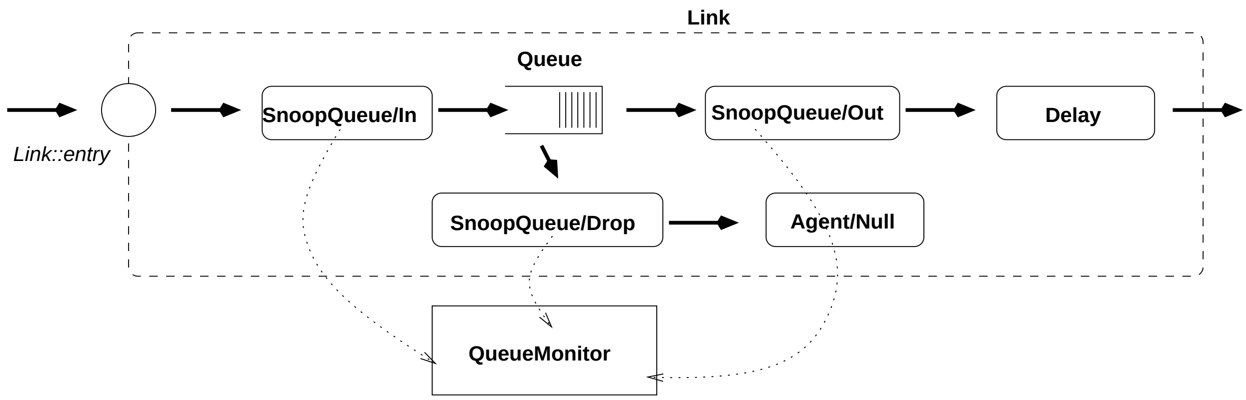

Figure 1: A QueueMonitor and supporting objects

Monitor Stats

- Simple stats kept by monitors:
  - arrivals (bytes and packets)
  - departures (bytes and packets)
  - drops (bytes and packets)
- Aggregate stats (optional):
  - queue occupancy integral
  - (bytes or packets)
- `QueueMonitor/ED` objects
  - "early" drops (bytes and packets)
  - some drops have this distinction (e.g. RED)
- Flow monitors:
  - types `QueueMonitor/ED/Flow` and `QueueMonitor/ED/Flowmon`
  - same as queue monitors, but also on per-flow basis
  - flow defined as combos of (src/dst/flowid)
  - flow mon aggregates and creates new flow objects

Break...

Emulation

- allows the simulator to interact with a real network (currently experimental and under development)
- Uses:
  - subject real-world implementations to simulated topologies
  - subject simulated algorithms to real-world traffic
- Traffic transducers:
  - real packets mapped to/from simulated packets via *tap agent*
  - real from either live wire or tcpdump-formatted file via libpcap
- Real-time scheduler
  - special version of (currently List-based) scheduler
  - ties simulated time to real-time
  - for now, punts if simulation gets far behind
  - (can still do interesting things!)
- Other helpful agents
  - icmp generation
  - ping responder
  - arp responder
  - NAT function

Tap Agents

- Maps simulated packets to/from real-world packets
- Associated `Network` object provides packet source/sink
- Packets are assumed to be network layer, generally IP
- Real traffic is stored in simulated packet's data area:

```
void TapAgent::recvpkt() // net->simulator
{
        ...
        // allocate packet and a data payload
        Packet* p = allocpkt(maxpkt_);
        // fill up payload
        sockaddr addr;  // not really used (yet)
        double tstamp;
        int cc = net_->recv(p->accessdata(), maxpkt_, addr, tstamp);
        // inject into simulator
        hdr_cmn* ch = (hdr_cmn*)p->access(off_cmn_);
        ch->size() = cc;
        ...
        double when = tstamp - now();
        ...
        if (when > 0.0) {
                ch->timestamp() = when;
                Scheduler::instance().schedule(target_, p, when);
        } else {
                ch->timestamp() = now();
                target_->recv(p);
        }
}
```

Network Object

- abstraction of a real-world traffic source/sink
- opened read, write, or read/write by caller
- base class for specific network types (e.g. IP network)
- Network class:
  - requires socket system API (UNIX or WinSock)
  - supports a basic send/recv interface
  - separate send/recv "channels" (i.e. sockets)
  - non-blocking optional
  - framework supports multicast, addr/iface selection, etc
- IP Network (`Network/IP` class)
  - RAW IP interface (requires privs for raw read/write)
  - multicast group join/leave
  - loopback on/off
  - multicast ttl
- UDP/IP Network (`Network/IP/UDP` class)
  - access to IP/UDP through underlying socket layer
  - does not require privs for use
- Frame Level Packet Filtering and Generation

Packet Capture Network Object

- Use LBNL's `libpcap` facility to access packet traces
- Provides access to trace files and live network packets
- common capabilities:
  - stats (packets captured and dropped)
  - packet filter function (using pcap bpf optimizer)
- Live Packet Capture/Generation (class `Network/Pcap/Live`)
  - promiscuous or ordinary packet capture
  - frame-level packet generation
  - specification of network interface name
- File Packet Capture (class `Network/Pcap/File`)
  - packet capture and filtering only

ARP Module

- not really an *agent* (not derived from `Agent` class)
- uses attached `Network` object for I/O
- provides ARP query/response processing
- also provides proxy ARP
- only works for Ethernet now
- (responder is currently under development)
- useful for generating link-layer headers
- ARP cache size configurable by user
- Methods:
  - `network` – set or gets associated Network object
  - `myether` – sets local ethernet address
  - `myip` – sets local IP address
  - `lookup` – looks up mapping for IP address
  - `resolve` – sends ARP request for IP address
  - `insert` – insert an entry into the ARP cache

ICMP Agent

- generates real-world ICMP messages
- class `Agent/IcmpAgent`
- currently supports only ICMP redirect generation
- Interface:

```
set become <who should have sent the redirect>
set target <who to send redirect to>
set dest <redirect for what destination>
set gw <new gateway to use>
$ia send redirect $become $target $dest $gw
```

- provides ability to masquerade as current default gateway (some systems check and require this to process a redirect)
- generates dummy packet (uses defunct GGP protocol) (inspected by victim host to determine which destination modify)

NAT Agent

- supports Network Address Translation (NAT) for use w/emulation
- Currently supports only TCP (UDP is an easy extension)
- Supported modes:
  - source address rewrite
  - destination address rewrite
  - source and destination address rewrite
- does not rewrite port numbers or ACK/sequence numbers

Open Discussion...

TCP Research with NS

- NS was developed largely for TCP research
- Studies have included:
  - congestion control behaviors (NewReno, SACK, ECN)
  - active buffer management (RED)
  - synchronization effects
  - unfairness studies/scheduling
- Industrial impact:
  - SACK is proposed IETF standard (RFC 2018)
  - ECN is experimental (RFC 2481)
  - RED is implemented in various commercial products
  - NewReno CC also experimental (RFC 2582)

TCP Tahoe

- First “modern” TCP, included with Berkeley Tahoe UNIX release
- Features:
  - fast retransmit
  - dynamic window sizing on congestion
  - exponential timer backoff and RTT estimation
- What is does:
  - slow-start on any loss
  - not ideal for large bw-delay product networks
- Demo...

TCP Reno

- Evolution of TCP tahoe
- Features:
  - fast retransmit
  - fast recovery
    - extra window "credit" for each received ACK
    - helps to "keep pipe full"
- What is does:
  - slow start on timer expiration
  - congestion avoidance on fast RTX with half-size window
  - big problem: partial ACKs can *deflate* window
  - window deflation kill ACK clock, requires timer expiration
- Demo...

TCP SACK

- Evolution of TCP Generally to Selective Repeat Protocol
- Features:
  - provide for selective repeat and selective ACKs
  - allows for better performance on lossy links
- What is does:
  - ACKs carry additional info in option area:
    * beginning and end of each block of info at receiver
    * used by sender to “fill in the holes”
  - regular ACK field has the final say
- Demo...

TCP ECN

- More drastic evolution (both endpoints and routers)
- Developed mostly for TCP so far, but applicable to others
- Features:
  - routers may provide info about their congestion state
  - endpoints may adapt (slow down) prior to packet drops
  - generally based on a queue *average* (e.g. RED)
- What is does:
  - sender sets “I am ECN capable bit”
  - routers set ECN bit on each packet when they are congested
  - ECN capable receivers (e.g. TCP) return ECN bit to sender
  - sender reacts to ECN bit similarly to a packet loss
  - sender indicates its reaction to its peer receiver
- Demo...

TCP and fairness

- TCP is well-behaved:
  - it uses end-to-end congestion control
  - it slows down when it loses packets
- Other protocols may not be:
  - UDP has no built-in loss detection
  - TCP will give up its bandwidth to ill-behaved flows
- Demo...

NS TCP Exercise 1

- Construct an ill-behaved TCP endpoint
  - it should still do some CC
  - make it “more aggressive”
- Run this TCP against a regular TCP
  - describe what happens
  - show it
  - how does it compare with an aggressive UDP flow?
- Some steps to accomplish this:
  - get into the ns source directory and create a new agent
  - make the simulator
  - create a topology and scenario to run
    - looking in `tcl/test` may be helpful

NS TCP Exercise 2

- Given an unfair set of traffic sources
  - construct a simulation using a "fair" form of packet scheduling
  - show this behavior using nam
- Some steps to accomplish this:
  - you will need to select some form of "fair" scheduling
  - look around for different Queue types