Proceedings of the 1989 Winter Simulation Conference
E.A. MacNair, K.J. Musselman, P. Heidelberger (eds.)

COMMUNICATION NETWORK ANALYSIS WITH COMNET II.5

Robb Mills
CACI Products Company
3344 North Torrey Pines Court
La Jolla, California 92037

ABSTRACT

COMNET II.5 is a general SIMSCRIPT II.5-based model for the performance analysis of circuit, message, or packet-switching networks. A variety of routing algorithms are built-in. After creating a network description with a convenient menu-driven editor, simulation follows immediately. There are no programming delays. Network operation is animated during the simulation. Reports show blocking probabilities, call queueing and packet delays, network throughput, link group utilization, and link group queue statistics. Some of the major enhancements that have been added to COMNET II.5 during the past year include capabilities for modeling multidestination messages, empirical probability distributions, flood routing, and access facilities. An access facility can be a point-to-point or multipoint line or a local area network.

INTRODUCTION

COMNET II.5 is a performance analysis tool for communication networks. Based on a description of a network and its routing algorithms, COMNET II.5 simulates the operation of the network and provides measures of network performance. No programming is required to use COMNET II.5. Network descriptions are created with a convenient menu-driven screen editor that is extremely easy to learn to use.

By using discrete event simulation methodology, COMNET II.5 provides realistic results. The drastic simplifying assumptions required when using analytical methods to analyze complex communication networks are eliminated.

New capabilities are added regularly to COMNET II.5 in order to stay current with the latest networking performance issues. The identification of potential new capabilities depends primarily on feedback from COMNET II.5 users. Many of the enhancements added during the past year have been aimed at fine-tuning the COMNET II.5 networking model. These enhancements are described in the next section.

COMNET II.5 can be used to simulate any communication network which uses circuit switching or store-and-forward operation. The network nodes in a COMNET II.5 model can represent circuit switches (PBXs, Central Office Switches, Time-Division Multiplexors), store-and-forward switches (packet switches, PADs, message switches, statistical multiplexors, concentrators, front-end processors, cluster controllers, routers, gateways), or a combination of both circuit and packet switching. Packet-switching networks with virtual circuit or datagram operation can be modeled. In the COMNET II.5 model, packet switching is used very generally to refer to store-and-forward operation. Access facilities connect to the network nodes in a COMNET II.5 model and can represent point-to-point or multipoint lines or local area networks. A variety of routing algorithms, including adaptive shortest path routing, are built-in.

The COMNET II.5 package consists of two parts: COMNETIN and COMNET. COMNETIN is the menu-driven, screen editor that is used to create and modify a network description. COMNET is the communication network simulation program. The primary input to COMNET is the network description created by COMNETIN.

COMNET prompts you for the name of the file created by COMNETIN and several simulation control parameters, such as the length of the simulation. After reading the network description file, the simulation begins. If you request animation of the simulation, you see a picture of the network in operation. Routing choices and changing levels of circuit group utilization are apparent.

COMNET II.5 is analogous in many ways to NETWORK II.5, which is another CACI product that is used for simulating the performance of computer systems and local area networks. NETWORK II.5 models processing, storage, and communication resources, while COMNET II.5 focuses primarily on communication resources. NETWORK II.5 facilitates modeling of general hardware-software interactions; COMNET II.5 facilitates modeling of circuit and packet-switching operation.

COMNET II.5 and NETWORK II.5 use different models for data transport that make COMNET II.5 better suited for modeling wide area networks and NETWORK II.5 better suited for modeling local area networks. COMNET II.5 uses a network level model of data transport, where the focus is on end-to-end and link group performance measures, as well as on routing issues. NETWORK II.5 uses a link-level model of data transport, where the focus is on individual link performance measures, as well as on link access protocols.

ENHANCEMENTS DURING THE PAST YEAR

Readers not already familiar with COMNET II.5 may want to skip this section and go directly to the following section on creating a network description with COMNET II.5.

Generalized Data Message Model

The data message model has been generalized so that it is now possible to model multidestination messages and triggered messages. A triggered message originates after the triggering message reaches its destination, typically following a user-specified triggering delay. One application of triggered messages is to model response messages that return to the origin of the triggering message. It is also possible to employ a user-provided file of data message traffic. In the generalized message model, message categories are defined in terms of origin, destination, and class of service, so that the message interarrival time distribution can now vary by class of service. The end-to-end acknowledgment frequency can be specified for each message category.

Enhanced Virtual Call Model

The response message size for virtual calls can now be modelled as a random variable, instead of a constant value.

A count of blocked virtual call attempts has been added to the Virtual Call Report.

Two new route selection criteria have been added for the user-defined node-by-node routing strategy for packets. According to one of the new criteria, a virtual call request packet is routed over the link group that has the smallest number of virtual calls in progress. Following the other new criterion, a virtual call request packet is routed over the first operational link group that has not reached its virtual call limit.

Enhanced Modeling of Random Variables

Users can define their own probability distributions by creating a table that lists each possible value of a random variable and its probability. It is also possible to model random variables with a piecewise linear cumulative distribution function (CDF). For a piecewise linear CDF, the table entries define the endpoints of the linear segments of the CDF.

The geometric distribution has been added to the list of distributions available for modeling discrete random variables, such as message sizes.

Up to 99 (increased from 9) random number streams are available for generating independent sequences of random numbers. Each random number stream corresponds to a different seed to the random number generator in COMNET II.5.

Expanded Packet Switching Model

Flood routing can now be modelled. With flood routing, a copy of a packet is routed over all outgoing links at a node, provided that the packet was not already flooded from the node and provided that the packet has not been circulating in the network longer than a user-specified timeout interval.

A load-balancing route selection criterion is now available for use with user-defined node-by-node routing tables. According to the load-balancing criterion, whenever multiple link groups are listed in the user-defined routing tables, a link group is selected at random from those link groups that are operational.

In previous releases of COMNET II.5, control packets (e.g., acknowledgments) have automatically had higher priority than data packets. Users can now override the default priority assigned to control packets.

In previous releases, the packet processing time did not vary with packet size. It is now possible for packet processing times to increase linearly with the size of the packet.

Access Facilities

An access facility is a new building block for a COMNET II.5 model. The purpose of an access facility is to model the contention delay that traffic encounters in moving between access devices (e.g., terminals) and a backbone network. Access facilities can be used to model polled multipoint lines, token ring LANs, or CSMA/CD LANS. A typical application is in predicting end-to-end delays for traffic that traverses a LAN before getting routed over a backbone wide-area network before crossing another LAN to reach its destination.

CREATING A NETWORK DESCRIPTION

A network description created by COMNETIN consists of three major categories of data: network topology, network traffic, and network operation.

Network Topology

The network topology is defined by the nodes of the network, the link groups interconnecting the nodes, and the access facilities assigned to each node. The interconnected nodes form the backbone of the network. The access facilities are used to get traffic to and from the backbone.

Nodes perform routing or switching functions and can also represent sources and destinations for network traffic. As mentioned in the introduction, nodes are used to model circuit switches (PBXs, Central Office Switches, Time-Division Multiplexors), store-and-forward switches (packet switches, message switches, statistical multiplexors, concentrators, front-end processors, cluster controllers, routers, gateways), or both.

A link (or circuit) group is a collection of identical, point-to-point transmission channels connecting exactly two nodes. Voice, data, and other types of traffic can be routed over the same link group. Multiple link groups can connect the same pair of nodes. Any backbone network topology in which a link group connects exactly two nodes can be represented.

An access facility connects to a single network node. Many access facilities can connect to the same network node. An access facility can be used to model a point-to-point or multipoint line, a token ring LAN, or a CSMA/CD LAN.

Node attributes include parameters such as the circuit-switched call setup time, the number of packet-switching processors per node, the packet-switching time per processor, and the total packet buffer space at the node.

Link group attributes include parameters such as the nodes at each endpoint, the number of links, the call priority required to access the group, an indicator if the group can carry tandem calls, an indicator if the group supports call queuing, the effective speed (bit rate) of each link, an indicator of half duplex or full duplex operation, and the propagation delay. The capacity of a link group can be defined in terms of the number of links (i.e., transmission channels) or the total bandwidth available. Link group availability is specified by the mean time to failure and the mean time to repair. Link group failures can occur at fixed intervals or at random; the repair time can also be constant or random.

Access facility attributes include the network access node, the number of devices contending for access, the propagation delay, the contention resolution protocol, and protocol-dependent parameters.

Network Traffic

There are three kinds of traffic that can be defined:

circuit-switched calls, data messages (which can be circuit-switched or packet-switched), and virtual-circuit calls. For a typical network, there are many categories of calls or messages. Each category is defined by an originating node, a destination node, and a class of service. Categories can also be defined in terms of originating access facility, destination access facility, and class of service.

The class of service can be used to define several categories of calls and/or messages with the same origin and destination, but with other characteristics (e.g., message size distribution or call holding time distribution) that vary. For circuit-switched calls, each class of service has a priority level and may have a bandwidth requirement and a call retry interval. The priority level determines which link groups a call is allowed to access and, for networks with preemption, which classes of calls can be preempted. The bandwidth requirement is used during routing to determine if a link group has enough unused bandwidth to handle a call. Blocked calls belonging to a class of service with a retry interval are reattempted after a fixed or random retry interval. Class of service is also used by some of COMNET II.5's algorithms for routing messages and the packets formed from messages.

In addition to the origin, destination, and class of service, other attributes for a circuit-switched call category include the call interarrival time distribution, the call holding time distribution, and an indicator if calls in the category can queue.

In addition to the origin, destination, and class of service, other attributes for each data message category include the message interarrival time distribution, the message size distribution, and, if applicable, the triggered message destination and class of service and the triggering delay. The destination can be given as a list of multiple destinations, causing a copy of the message to be broadcast to all destinations in the list. Depending on the message size distribution, data messages can represent batches of data ranging in size from short inquiries or requests up to large files. If a data message is circuit-switched, the holding time for the message is determined by the message size and the speed of the circuit-switched connection. If a data message is packet-switched, the data message is broken into packets and the packets are transported through the network using datagram operation.

Virtual-circuit call traffic is used for the traffic load in packet-switched networks which transport packets using virtual-circuit operation. Additional attributes for each virtual-circuit call category include the call interarrival time distribution, the messages-per-call distribution, the message interarrival time distribution, the message size distribution, the probability of a response message, the response message size distribution, and the delay before transmission of a response message.

The attributes of many of the traffic categories are random variables with some probability distribution. An extensive list of distributions is available in COMNET II.5 for modeling both discrete and continuous random variables. In addition, it is possible for users to define their own probability distributions for use in generating the traffic load on the network. For cases where actual or historical traffic data are available, it is possible to drive the simulation with files of user-provided traffic data, rather than relying on the generation of traffic data based on probability distributions.

Network Operation

Network operation parameters include a description of the network's routing algorithm. The routing algorithm determines how the network traffic moves through the network topology from origin to destination. COMNET II.5 includes both static and adaptive routing strategies.

With static routing, routing tables are predetermined. Node-by-node routing tables contain lists of link groups. The correct list is selected based on the kind of traffic, the node at which a routing decision is being made, and the destination node. A link group is selected from the list according to the type of traffic and a user-specified selection criterion. A special type of static routing is flood routing, where a packet is routed over all outgoing link groups. Source-node routing tables contain lists of end-to-end paths. Each path is a sequence of link groups from source node to destination node. There is a separate table for each category of traffic.

With adaptive routing, the routing tables are updated dynamically during the simulation so that traffic is routed along the shortest path to the destination node. The shortest path depends on the distance metric used for the link groups in the network. One example of a distance metric is the average link delay during some measurement interval. In this case, the shortest path through the network corresponds to the path with the shortest expected delay. Another type of distance metric allows the user to assign a set of weights or penalties to a link group. As traffic on the group increases, the length of the link group increases to the next penalty level.

Examples of additional network operation parameters include packet sizes, protocol overhead levels, retransmission intervals for blocked packets, the measurement interval for updating the shortest path calculations, and end-to-end flow control methods for virtual calls.

THE COMNET SIMULATION MODEL

Circuit Switching Model

For circuit-switched calls, COMNET models two different routing strategies: node-by-node routing and source-node routing.

With node-by-node routing, each pair of nodes has a user-defined list of link groups. Each ordered list contains the link group choices for routing a call to the next node, given the current node and destination node.

Each circuit-switched call that is generated during the simulation proceeds through a connection setup process. The setup process begins at the source node and proceeds node by node until the destination node, D, is reached or the call is blocked. When a call has reached some node J along the setup path, the routing table associated with the node pair (J,D) gives the ordered list of link group choices for moving to the next node on the setup path.

The call is routed over the first operational and accessible link group in the list that has sufficient idle capacity. The capacity of a link group can be defined in terms of number of channels or total bandwidth. If capacity is defined in terms of channels, each call requires one channel of capacity. If capacity is defined in terms of bandwidth, each call will require an amount of bandwidth determined by the call's class of service. A link group is considered accessible if a call's priority is at least equal to the access level of the link group and, for tandem calls, if the link group is permitted to carry tandem traffic. A call becomes a tandem call when it is no longer being routed from its source node. If an outgoing route is found, the necessary capacity is acquired on the chosen route and a signaling delay is incurred before making a routing decision at the next node.

If no outgoing route is found, but preemption is allowed, call setup may still proceed if a lower priority call or calls can be preempted. If an outgoing route is still not available, but the call is allowed to queue, the call is inserted in the queue of the first operational and accessible link group in the routing table that has room left in its queue. The call then waits its turn in queue until sufficient capacity becomes available on the circuit group, at which time call setup resumes.

Once call setup reaches the destination node, the capacity acquired along the connection path is not released until the holding time elapses (unless the call is preempted or disconnected by a link group failure).

If a node is reached from which no outgoing route can be found and preemption and queueing are not possible, the call is blocked. A blocked call is cleared or reattempted some time later, depending on the call category. In clearing a blocked call, any capacity acquired on a partially established connection path must be released.

Since there are typically signaling and sometimes queueing delays during connection setup, there are generally multiple connection setup processes active at the same time.

With source-node routing there is an ordered list of end-to-end paths for each call category. An end-to-end path is a sequence of link groups from the source to the destination node.

When a call originates during the simulation, it is routed over the first path in the list that has all link groups operational and sufficient idle capacity on each link group of the path. If none of the paths can carry the call, but preemption is allowed, the call may still be established by preempting a lower priority call or calls. Preemption is attempted beginning with the first path in the list and continuing until a path is found or the list is exhausted.

Once a path is found, the call acquires the necessary capacity on each link group. Unless the call is preempted or disconnected by a link group failure, the capacity is not released until the holding time for the call has elapsed.

If no path can be found, the call is blocked. A blocked call is cleared or reattempted some time later, depending on the call category.

In contrast with node-by-node routing, the routing and setup process for source-node routing occurs instantaneously in the COMNET II.5 model, which is equivalent to treating call setup and routing delays as negligible relative to the length of a call.

When a call termination event occurs for either a node-by-node or source-node routed call, the capacity used at each link group along the connection path is released immediately. In addition, a circuit group contention manager is executed in case any other traffic is waiting to use the link group capacity that just became available.

Packet Switching Model

In the COMNET II.5 model of store and forward operation (which is referred to here as packet switching), packets encounter, while being routed from source node to destination node, five components of end-to-end delay. Upon arrival at a node, a packet encounters (1) a processor queueing delay while the packet waits in queue to be processed or switched, (2) the actual processing or switching time during which a routing decision is being made in order to assign the packet to an outgoing link, and (3) a transmission queueing delay while the packet waits in queue to use the selected link. During actual transmission from one node to the next, the packet encounters the remaining components of delay: (4) the link transmission time and (5) the propagation delay.

When a data message originates during the simulation, the message is broken into packets (if necessary) and the packets are inserted in the processor queue at the originating node. In addition all packets arriving at a node on incoming links from other nodes are inserted in the processor queue. Packets are ordered by packet priority in the processor queue and the link group queues.

With virtual circuit operation, when a virtual call originates during the simulation, a virtual call request packet is inserted in the processor queue at the originating node. The call request packet is routed to its destination and a call connected packet is returned to the origin over the virtual circuit path established by the call request packet. Upon return of the call connected packet, messages are generated according to the message interarrival time distribution for the virtual call category. The messages generated during a virtual call are also broken into packets, but instead of being routed independently, each packet is constrained to follow the virtual circuit path defined by the call request packet.

Each packet switching node has one or more processors operating in parallel to remove packets from the input queue. After a packet is removed from queue, it incurs a packet processing or switching time that depends on the type and size of the packet and the packet switching speed. After the processing or switching time, an outgoing link group is selected for the packet (unless the packet has already reached its destination). The route selection logic depends on the type of packet. For example, a datagram packet is routed independently, whereas a virtual-call information packet must follow the virtual-call connection path already established by a virtual-call setup packet. If the outgoing route has an idle link, the link is made busy and a transit-packet arrival is scheduled at the next node after the transport delay for the packet. Otherwise, the packet is added to the queue for the outgoing link group. Packets blocked, perhaps due to insufficient buffer space or failure of all possible routing choices, are retransmitted after a user-specified timeout interval.

Upon execution of a transit-packet arrival event, the incoming circuit is idled and COMNET's circuit group contention manager is executed in case any other traffic is waiting to use the link group. The circuit group contention manager gives highest priority to circuit-switched calls; otherwise, as a general rule, traffic is handled on a first-come, first-served basis. After any propagation delay passes, the arriving packet is inserted in the queue for the packet switch at the receiving node.

Whenever a link group failure event occurs during the simulation, routing tables are updated to reflect the failure. All calls using the failed link group are disconnected and all virtual calls are rerouted. A link group restoration event is then scheduled. When the restoration event occurs, the routing tables are again updated and the next failure event is scheduled for the link group.

NETWORK PERFORMANCE MEASURES

COMNET II.5 produces several reports summarizing network utilization and network service. The Circuit Group Report for Circuit-Switched Traffic shows circuit group utilization and call and message statistics by circuit group. The statistics include blocking probabilities and queueing delays. The Circuit Group Report for Packet-Switched Traffic shows circuit group utilization, buffer use, and packet statistics by circuit group. The Packet Switch Report shows buffer utilization, packet switching delays, and packet throughput statistics. The Access Facility Report shows utilization, throughput, and delay statistics for each access facility. The Circuit-Switched Call and Virtual Call Reports show the network service received by each category of call. An example of a measure of network service for a circuit-switched call is the end-to-end blocking probability. The Message Report shows the network service received by each category of message. Examples of measures of network service for a message category are the average, standard deviation, and maximum response time. The Disconnected and Preempted Calls Report gives, for each call category, a count of the number of calls disconnected due to link failures and the number of calls preempted by higher priority calls. The Call Routing Report gives, for each call category, the number of calls routed over each possible end-to-end path defined in the source-node routing tables.

AUTHOR'S BIOGRAPHY

ROBB MILLS is the COMNET II.5 Product Engineering and Development Manager at the CACI Products Company. In addition to directing the development of COMNET II.5, he teaches regularly scheduled classes on network performance analysis using COMNET II.5 and consults with CACI clients on modeling and simulation projects. Prior to joining CACI in 1986, he was with AT&T Bell Laboratories for 12 years, where he completed various systems engineering and operations research assignments for telecommunication systems and networks.
He has bachelor's and master's degrees from Cornell University in industrial engineering and operations research and a doctorate from Columbia University in operations research.

CACI Products Company
3344 North Torrey Pines Court
La Jolla, CA 92037
(619) 457-9681