Infrastructure Optimizer
Modeling of Data Link Components

HyPerformix, Inc.
4301 Westbank Drive
Building A, Suite 300
Austin, TX 78746
www.hyperformix.com

Abstract
This whitepaper presents an overview of the HyPerformix Infrastructure Optimizer™ data link models.

Introduction

This document describes the following Optimizer network components: standard LAN components, such as Ethernet, Token-ring and FDDI; WAN components, such as the Basic WAN and Frame Relay; and generic datalink components, like Datalink and Point-to-Point. This paper also describes how to use the components in common data networking scenarios.

Each of the components has a set of parameters associated with it. All of them have name, configuration, cost and performance parameters. The name of each component must be unique within a model. The cost is used to represent a monetary value and is only used in model documentation. The configuration provides a drop-down list of pre-defined performance characteristics of commonly used products, or the user can define new configurations. The performance parameters may also be set individually by the user.

Standard LAN Components

Optimizer provides three standard LAN components: Ethernet, token ring and FDDI. They are configured with two parameters: bandwidth and background load. Since these types of LANs have been around for a long time, their performance characteristics are well-known. Optimizer performance models of standard LAN components take advantage of the knowledge that has been gathered during measurement studies of the behavior of standard LANs.

A packet transferred across a LAN is encapsulated in a LAN frame. During simulation, Optimizer determines the queuing delay suffered by a LAN frame based on the LAN traffic at the time of arrival of the frame. In addition, a fixed 10 microsecond access delay to model NIC setup, inter-frame gaps, and arbitration resolution is associated with each frame. Finally, the time due to data transmission is computed according to the following formula:

```
Transmission Time = L / (s * (1 - b)),
```

where L is the length of the LAN frame in bits, s is the LAN bandwidth in bits per second, and b is the background load as a fraction. Standard LAN frame overheads accounting for frame header and trailer are used by Optimizer in determining the length L of a frame.

Basic WAN Components

Basic WAN

The basic WAN component is useful for modeling networks about which very little is known; e. g., the Internet or a frame relay core network. The basic WAN component has only one performance parameter: latency. The latency parameter specifies the delay a packet suffers when it is transferred across the WAN. For the WAN component, the latency value includes all queuing delay, propagation delay and transmission delay suffered in the WAN. In such networks, the values for queuing delay and transmission delay are often not available, but an approximate value for the total delay may be specified by the service provider. This total delay value can be used as the latency in Optimizer configuration.

Frame Relay

The Frame Relay component is a specialized Basic WAN component. One would use a Frame Relay to specify one or more frame relay permanent virtual circuits (PVC). This component must be connected to a Point-to-Point and Router components on both sides to form a Router – Point-to-Point – Frame Relay – Point-to-Point – Router sequence. Several Router – Point-to-Point pairs may be connected to the same Frame Relay component (i.e. to create a star formation). Within the frame relay component, PVCs are defined from one Router to another Router to define a data path through the Frame Relay. The committed information rate (CIR) and latency are specified for each PVC. This latency parameter represents the total delay through the PVC.

Several throughput statistics are reported for Frame Relay: total throughput and data throughput, interval throughput (based on the user-specified CIR computation interval parameter), and excess throughput (which is in excess of the user-specified CIR).

Generic Datalink components

Generic Datalink

The Optimizer datalink component is a generalized link that is suitable for modeling many network components which are not represented with a specific icon in Optimizer.

The Datalink component is configured with the following performance parameters:

bandwidth
latency
arbitration time
maximum payload bytes
minimum payload bytes
overhead bytes
simultaneous streams
background load

A data packet transferred across a Datalink is encapsulated in a frame. The format of the frame, that is: maximum payload, minimum payload and overhead bytes, is specified by the OSI layer 2 protocol used to transfer packets across the Datalink. The data packet is included in the frame as the payload. The size of the header and trailer of the frame equals the overhead bytes.

Each frame going across a Datalink will suffer a delay consisting of four components: queuing delay, arbitration time, transmission time, and latency.

The queuing delay depends on how many frames are waiting for transmission and the remaining transmission time of any frame that is being transmitted at the time a frame arrives at a Datalink. Some frames may suffer no queuing delay and some frames may suffer a large queuing delay. In Optimizer, the queuing delay of a frame is determined internally during simulation.

The arbitration time and latency are fixed values. Each frame suffers the same delay due to arbitration needed to gain access to the Datalink for transmission. In most cases the arbitration time will be negligible compared to the transmission time. The latency is equivalent to the propagation delay of a

frame on the Datalink, and any fixed delay due to network overhead. For short links, the propagation delay is negligible. For long links (e.g., satellite links), it may be significant.

Optimizer uses the bandwidth value and the background load value to compute the transmission time of a frame across the link. If the Datalink bandwidth is s bits/sec and the background load is b (a fraction), then the transmission time of a frame of length L bits is computed as

```
Transmission Time =   L / (s * (1 - b)) seconds.
```

In Optimizer, background load is specified as a percent. Internally it is converted to a fraction.

The simultaneous streams value is used to specify how many streams of data can be transmitted over the Datalink simultaneously. If the simultaneous streams value is set to 2, then the link can be used to send frames simultaneously in the two directions. Thus, the link may be full-duplex. The bandwidth value specified in the configuration is available in both directions along the same data path or along one direction at a time for two data paths – as demand requires. If the simultaneous streams value is set to 1, then the link is shared (i.e., half-duplex); only a single packet can be sent in any direction at a time. If four components are connected to a Datalink with four simultaneous streams, the Datalink can support full duplex data transmission along two independent data paths. E.g., simultaneous full duplex data transmission could be supported between A-B and C-D. However, data transmission between A-B and C-B might compete since B can only support one input and one output stream of data at a time.

Point-to-Point

A Point-to-Point component is a specialized generic Datalink component. Point-to-Point is slightly more restrictive than generic Datalink in their allowable topology connections and configuration. Point-to-Point components may only be connected to a single interconnect, frame relay, WAN or computer on each endpoint. Their configuration parameters include half and full-duplex as compared to the Datalink's simultaneous streams parameter, and rely on the default arbitration delay of 10 microseconds as compared with the Datalink's user specified arbitration delay. A full-duplex Point-to-Point may be configured with asymmetric performance characteristics, including bandwidth. In addition, a full-duplex Point-to-Point will report separate statistics for data flow in each direction.

Simulation of Network Interface Cards

Optimizer contains a generalized network interface card (NIC) model. A fixed 10 microsecond access delay is used to model NIC setup, inter-frame gaps, and arbitration resolution. This delay was discussed earlier as the arbitration delay. Several algorithms are employed by Optimizer to model NICs that are contending for access to a datalink. Token rings and FDDI devices use a round-robin algorithm to select a NIC. Once the NIC is selected, it retains access to the datalink for several packet transfers using a limited-k scheme. Use of limited-k to approximate target token rotation timer (TTRT) behavior greatly reduces model complexity and improves simulation time. This scheme is long enough (in terms of bits) to provide each NIC with exhaustive service. The round-robin ordering is set according to an internal NIC numbering scheme. Each NIC has its own unique number, the numbers correspond to the topology diagram – and may be influenced by the order in which the topology diagram is drawn. The effect of this on end-to-end response time, however, is felt to be negligible.

Ethernet NICs, however, use sampling from a triangular distribution to grant access to NICs. This model exhibits the correct behavior in the common case (i.e., when there is no contention and subsequent collision). Moreover, it approximates the capture effect that enables a NIC to transmit a number of

successive frames that takes place during periods of high contention, which is expected to occur in high performance/high traffic servers.[1]

Some Applications of Data Link Components

In this section we describe the parameter values for the Optimizer datalink component in standard data networking applications.

Bandwidth

Kilobits / Megabits

Optimizer GUI provides a list consisting of three scaling factors for specification of bandwidth: Bits/sec, Kilobits/sec, Megabits/sec. As explained in the Optimizer online documentation, kilo is equal to 1024 and mega is equal to 1048576 in the context of specifying bandwidth. In the transmission time formulas presented earlier, 's' is the equivalent in bits/sec of the number specified for bandwidth during configuration of the datalink. Thus, an Ethernet with bandwidth specified as 10 Mbits/sec in the GUI, will actually have a bandwidth of 10485760 bits/second. Hence, in order to specify a bandwidth of 10000000 bits/second, it is necessary to configure the Ethernet bandwidth as 10000000 Bits/second in the GUI. This observation about the usage of mega applies to FDDI, Token Ring, generic Datalink, and Point-to-Point as well.

Overhead

The circuit bandwidth of some Point-to-point configurations is higher than the bandwidth available for a layer 2 frame transmission. This is because part of the circuit bandwidth is used as overhead. In the following table, note that link bandwidth is the bandwidth available to transmit a frame.

Point-to-Point Link Name	***Circuit Bandwidth (bits/sec)***	***Link Bandwidth (bits/sec)***
Dataphone Digital Service (DDS)[2]	64000	56000
ISDN BRI (Basic Rate Interface)[3]	144000	128000
T1[4]	1544000	1536000
T1 D4 AMI[5]	1544000	1344000
Fractional T1[6]	1544000	64000 - 1472000

[1] This behavior is described in Rich Seifert's, "The Effect of Ethernet Behavior on Networks Using High-Performance Workstations and Servers", March, 1995, Technical Report, Networks and Communications Consulting, Cupertino, CA.

[2] One of every eight bits transmitted in the circuit is used for synchronization between the transmitting and receiving ends.

[3] Channel configuration: 2B + D = 2*64000 + 16000; where B carries data and D is for signaling.

[4] 8000 time slots are present every second, and 24 channels are multiplexed together. Each channel transmits 8 data bits plus 1 bit to represent a time slot boundary. From Feibel's "Complete Encyclopedia of Networking", the samples for the 24 channels yield 1.536 Mbps, and 8 kbps is added for framing to make the 1.544 Mbps total circuit capacity for a T1 line.

[5] With the older T1 technology (D4 AMI), the subscriber only gets 56 kbps of every channel; the service provider steals one bit of each value for control purposes. 56 Kbps times 24 channels yields 1.344 Mbps.

[6] Consists of 1 to 23 T1 channels = 1 * 64 Kbps to 23 * 64 Kbps = 64000 to 1472000 bits/sec.

Link layer Protocol: PPP (Point-to-Point Protocol)

Data networking traffic is forwarded across a wide area link using a link layer protocol. There are several protocols used in practice. The following is a list of Optimizer parameters for a common protocol.

PPP may be used over a full-duplex link between two points. A router supporting PPP over a datalink will let the user configure the maximum payload bytes. Often it is set to the same value as the Ethernet maximum payload bytes (1500 bytes) if the router connects to an Ethernet and will forward packets to and from the Ethernet. This information is exchanged between the two ends of a PPP link with link control protocol packets which are used for hand shaking between the two ends. Each PPP frame has a frame header and a frame trailer. The total size of them is ten bytes. Thus to model a PPP link, a data link component should have the overhead bytes set to 10.

PPP can be configured to support data compression. If you are using this feature, then the effect of data compression can be taken into account in the model by increasing the bandwidth by a factor equivalent to the extent of data compression.

Further Information

Two other documents that supplement this article are available. They give details of how the network and interconnect devices are modeled in Optimizer:

1. Infrastructure Optimizer Modeling of Networks
   http://www.hyperformix.com/whitepapers/modeling-networks-whitepaper.pdf
2. Infrastructure Optimizer Modeling of Interconnect Components
   http://www.hyperformix.com/whitepapers/modeling-interconnects-whitepaper.pdf

