Performance and Capacity Analysis of an Operating ATCS Communications System

EDWARD L. FURMAN

Automated Monitoring and Control International has been working with Union Pacific Railroad to implement a communication network based on Advanced Train Control Systems (ATCS) Specification 200. This is the first large-scale implementation of the ATCS communications systems. Computer performance predictions and data from the installed system give insight into the capabilities of this type of mobile data network. Expected throughput of the network is estimated on the basis of analytical models. The successful large-scale implementation of a Specification 200 network on Union Pacific indicates that the ATCS specifications provide the basis for a viable communication network.

The use of mobile data technology is becoming more important in the transportation industry. In the railroad industry, business and safety systems have extended the reach of host computer networks and dispatch operations to the locomotive through the use of data radio communications. Business applications such as Automated Monitoring and Control International's (AMCI's) Work Order Reporting System (WORS) give railroads a distinct competitive advantage through improved productivity and customer service. WORS is designed to operate within the architecture established by North American railroads for the Advanced Train Control System (ATCS) (*1*).

Figure 1 shows the components required to complete an end-to-end data connection in ATCS. The dispatch computer determines an appropriate movement authority and creates ATCS messages to the locomotive. This movement authority is delivered and displayed to the engineman in the locomotive. The communication process begins as the dispatch computer delivers the message to the front-end processor (FEP), which then sends the message to the appropriate cluster controller (CC). The CC selects the optimum base communication package (BCP) to deliver the message to a mobile communication package (MCP) located on the locomotive. The BCP transmits the message to the MCP over a 900 MHz radio frequency (RF) path. When the message is received by the MCP it is then routed to the on-board computer (OBC), which drives a display that shows the movement authority to the engineman.

AMCI has been working with Union Pacific Railroad to implement a communication system based on ATCS Specification 200-compliant components. A locomotive in WORS communicates over the same network as an ATCS locomotive, but the messages are sent between a conductor's terminal on the locomotive and the railroad's management information system (MIS) host computer system. The FEP routes WORS traffic to the MIS host instead of the dispatch computer. On a locomotive in WORS the MCP connects directly to an on-board terminal (OBT) instead of the on-board computer. With the addition of such an on-board computer, a WORS locomotive can be expanded to full ATCS.

WORS is a system that enables crewmen to report car movement information directly into the railroad's central computer system from terminals on board the locomotives. Work is reported in a real-time manner rather than through after-the-fact clerical input. Train operations, such as train movement data, delivery and collection of cars, car status updates, and interchange of cars with connecting carriers, are entered at the conductor's on-board terminal. If the planned train work needs to be modified, the modifications are transmitted to the train. This approach avoids after-the-fact clerical input and manual data entry tasks, which in turn increases productivity and dramatically improves operational data reliability. A reduction in car-hire costs and improved customer service are additional benefits.

The communication network operates according to specifications established by the ATCS project of the Association of American Railroads (AAR) and Railway Association of Canada (RAC). These specifications are contained in the ATCS Specification 200 series. The RF link uses a pair of radio channels in the 900 MHz band. Each channel operates at a data rate of 4,800 bps. One channel (the inbound channel) receives data from the locomotives operating within the coverage of a base station. The second channel (the outbound channel) carries transmitted data from the base station to the locomotives.

Messages sent over the link are separated into individually addressed packets of data. Messages longer than 251 bytes require multiple packets for transmission. Two primary modes of packet transmission are used. Mode 1 requires an acknowledgment packet upon receipt of each packet. Mode 2 uses an acknowledgment or resend request for groups of 10 packets rather than for each packet. The last two packets of each group of 10 request acknowledgment.

WORS is the first large-scale implementation of ATCS technology in a revenue-producing environment. In terms of how this Specification 200 communication network performs, such evaluations are generally made by considering throughput and response time. The WORS network at its current load is not operating near the expected capacity. Throughput and response time under heavily loaded conditions have only been experienced during laboratory testing and through computer simulation of the network. But it is possible to estimate the expected throughput of the network and to review the

Automated Monitoring and Control International, Inc., 11819 Miami Street, Omaha, Nebr. 68164.

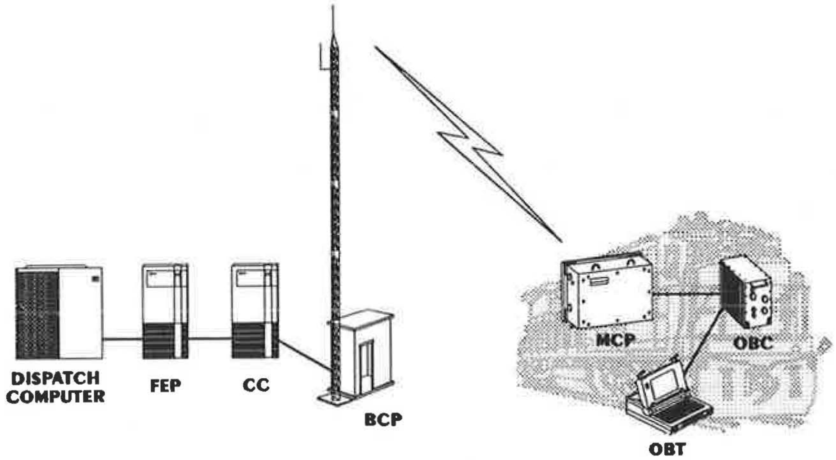

FIGURE 1 ATCS communication system.

current loading of the network and to examine mechanisms for adjusting the capacity of this type of network.

NETWORK CONFIGURATION

As of March 15, 1991, 67 base communication packages were installed at Union Pacific sites providing radio coverage in the states of Oregon, Idaho, Washington, Utah, Wyoming, and Nebraska. Figure 2 shows the Union Pacific Railroad and the areas currently using this ATCS-compliant data radio system. More than 500 locomotives are equipped with mobile communication packages. The communication network is controlled from a centralized fault-tolerant combination front-end processor/cluster controller. On average, 90 jobs use this network every day for WORS. The base communication packages are connected to the front-end processor/cluster controller on a point-to-point basis using the ATCS specification balanced HDLC protocol. Currently all the base communication packages are operating on ATCS Channel 1. The mobile communication packages are able to operate on any of the six ATCS channels.

The traffic currently on the network is of three types: basic operation of the network, WORS, and 3270 pass-through. All messages routed through the network conform to ATCS Specification 250. The messages result in packets of varying sizes. Outbound messages for WORS and 3270 pass-through utilize the extended ATCS packet length of 256 bytes. Most inbound messages in WORS and 3270 pass-through are less than 128 bytes and utilize the standard ATCS packet length of 128 bytes.

The number of locomotives using a base communication package in the operating WORS network was evaluated using the front-end processor/cluster controller's diagnostic function, "LIST DEVICE CONTACTS." This function produces a report for selected base communication packages showing a list of locomotive unit numbers and the time when each of those locomotives last used the base communication package. The reports are evaluated by considering a 1-hr window at and before the sample time and counting the number of locomotives reported during that time. This is based on the assumption that sometime during the hour all of these locomotives may have been using the base station concurrently. Based on these reports, as many as 29 locomotives are operating through a given base communication package.

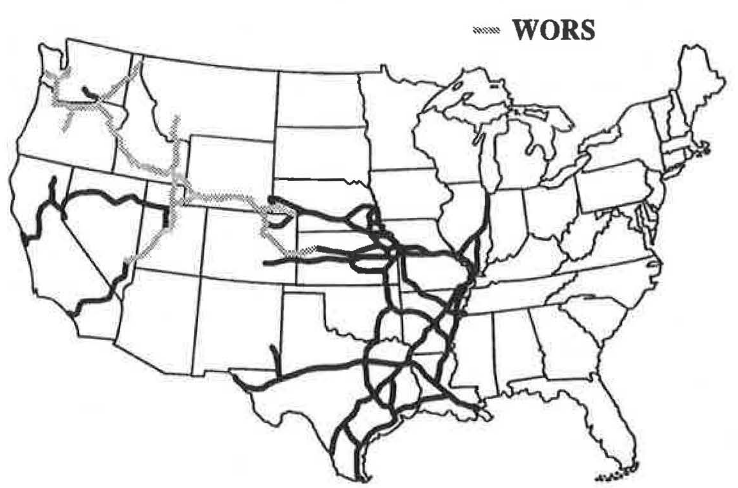

FIGURE 2 Union Pacific Railroad system map.

EXPECTED CAPACITY

RF Channel Capacity

The capacity or throughput capability of the network is primarily determined by the capacity of the RF link. The link must be further separated into the inbound and outbound directions because different performance considerations are involved for each direction of traffic. Capacity can be considered in packets or bytes. Because the available performance statistics from the front-end processor are in packets, most of the capacity discussion will be in terms of packets. This introduces a level of inaccuracy because the system uses packets ranging in size from 8 to 256 bytes. A lower limit for packet capacity can be estimated by using the largest packet size.

In the outbound direction, the capacity is determined by the data rate over the RF. Because a single base is used to reach multiple mobiles under the control of the front-end processor/cluster controller, it can be assumed that the base transmissions can be scheduled to utilize fully the available capacity of the base. Throughput will be estimated in packets per hour. For the largest packet size—256 bytes—approximately 5,000 packets can be sent from a base station in 1 hr. If smaller-size packets are transmitted then, of course, more packets per hour can be sent over the base.

The inbound channel is characterized by multiple mobiles trying to communicate with a single base. When multiple mobiles attempt to use the same channel more than one mobile may attempt to transmit at the same time. This can result in a "collision" and both mobiles will have to retransmit the packets. The inbound capacity is determined by the effectiveness of the mechanism used to minimize inbound collisions.

There are two possible modes of inbound channel operation, both of which are used in the network. One mode, further described below, is by "busy bit" control. The other mode, called ALOHA, is where mobiles simply access the channel and transmit a packet. If the packet fails, it is retried. From Figure 3, for $a = 1$, the inbound capacity of ALOHA for 256-byte packets is approximately 1,000 packets per hour.

The ATCS radio link uses digital sense multiple access (DSMA) to control inbound channel access (*2*). When a base station is transmitting data outbound, it will insert busy bits in the data stream upon detection of inbound traffic. The busy bits are monitored by mobiles wanting to use the channel.

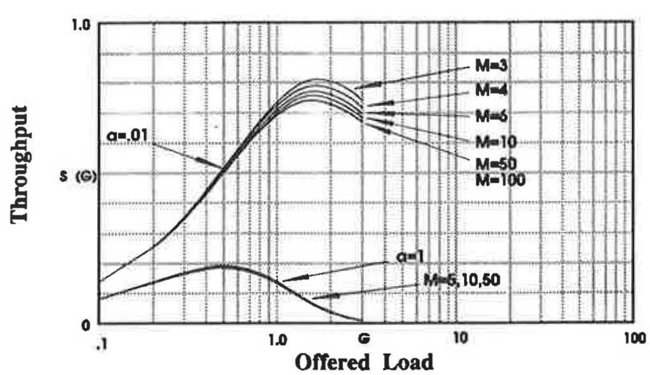

FIGURE 3 Normalized throughput for inbound channel.

When a mobile communication package checks the channel and finds that some or all of the last three busy bits were set to busy, it will wait a random time from 10 to 2,000 ms before checking the channel again. If the mobile finds the last three busy bits set to idle when it checks the channel, it will access the channel and transmit a packet. If a sync pattern is detected, the mobile will wait a random time from 10 to 800 ms and check the channel again. If less than three busy bits have been received since the sync pattern, the mobile will wait a random delay of 10 to 50 ms and check the channel again.

AMCI developed an analytical model to predict the capacity of the ATCS inbound channel (*3*). Figure 3 presents the performance prediction of this model in the form of normalized throughput curves. This graph considers throughput in the manner used by Kleinrock and Tobagi (*4*) and assumes that all the packets are the same length. This method of looking at the inbound channel removes the throughput differences related to packet length by normalizing all time to the time required to transmit a packet.

The y-axis is the throughput normalized by dividing the number of packets that got through by the time required to transmit one packet. The x-axis is the load presented to the channel in packets, normalized by dividing by the time required to transmit one packet. The throughput is a function of the parameter a, which represents the portion of time when each mobile has no information about other inbound traffic on the link. When $a = 1$, mobiles have no information about other traffic and access the inbound channel randomly. This results in the lower throughput curves in Figure 3, which represent the well-known performance of ALOHA channel access (*4*). For a carrier sense multiple access (CSMA) system, a is primarily a function of propagation delay. In the AMCI model for the DSMA inbound channel used by ATCS, the parameter a is determined by modem bit synchronization time and busy bit insertion time in addition to propagation delay and other factors. Figure 3 shows the changing throughput characteristics for various numbers of mobiles (M) for an assumed a of 0.01. The curves in Figure 3 show that the throughput of the channel increases as more load is offered to the channel. For very high offered loads the throughput will decrease as more collisions occur. The family of curves also shows that with fewer mobiles there are fewer collisions and greater throughput is possible.

For any given packet length the normalized throughput curve can be converted to a packet throughput for a specific time period. Assuming a 256-byte packet, a lower limit of inbound capacity of a base station is estimated from Figure 3 to be approximately 4,200 packets per hour.

Required Capacity for ATCS

The expected operating scenario for a railroad's ATCS implementation determines the characteristics of the required data traffic. The key parameters of this traffic are not only data volume but also the geographic location of the mobiles sending and receiving the traffic as well as the distribution of wayside and other devices in the system.

The RF channel requirements for ATCS were estimated by Dr. Sheikh of Lapp Hancock (memorandum from R. Ayers to Component Specification Drafting Committee members, Oct. 9, 1987). This requirement estimate was in the form of multiple curves produced for various situations. As an example from this estimate, three radio channels were sufficient to support the following load:

- 20 trains,
- Location reporting every 60 sec,
- 70 data base downloads per hour,
- 350 wayside units, and
- Status updates every 90 sec.

Many of the control flows and operating assumptions of ATCS are still under development, so the exact communication requirements are not known at this time.

RF Capacity Expansion

The amount of traffic that a base communication package receives is determined by the number of locomotives and other devices within its coverage area. The traffic load on the system further increases when base communication packages are located to provide overlapping coverage. Generally, the larger the coverage area of a base communication package, the more locomotives and other devices it will be communicating with. The coverage area of a base station can be controlled by antenna selection and mounting height.

When a base communication package is installed at a specific location with an antenna of known pattern at a given height, a coverage area can be calculated. Each base communication package has a known capacity for inbound and outbound traffic and will provide this capacity for all the trains, wayside units, and other RF users within its coverage area. If the traffic demand within the coverage area is expected to exceed the capacity of the base station, the network capacity can be expanded by installing additional base communication packages on other RF channels. Six channels have been licensed to the AAR for ATCS. If more capacity is required in an area than can be provided by base communication packages on six channels, the network capacity can be expanded by reducing the coverage area of each base station and installing additional base communication packages with smaller coverage areas.

Known values of transmitter power and receiver noise figure are used to compute a carrier-to-noise ratio (C/No) for RF data reception. AMCI uses a computer model in conjunction with digitized topographical data to calculate contours where the C/No is acceptable for a given level of service. This predicted coverage area identifies the physical locations where users can expect to obtain normal performance from the data radio network. Field tests have established the accuracy of the coverage prediction (*5*). Coverage analysis is used in conjunction with railroad operating plans to determine the traffic volume expected for a base station. Base station locations and antenna configurations are selected for acceptable system loading and to provide the required level of service to mobiles.

An ATCS communication network can be engineered to meet the capacity required by the railroad. The areas in coverage and the capacity of the RF network are controlled by engineering the coverage of base communication packages. The remaining components in the network can be expanded as needed to support the base stations installed.

SYSTEM PERFORMANCE

Total Traffic Volume

A macro view of the WORS ATCS Specification 200 network in operation is available from a statistical report produced by the front-end processor/cluster controller. This report counts the packets received and transmitted for each hour by each base communication package. Data from this report for March 15, 1991, a randomly selected day, were used to graph the inbound and outbound traffic per base station per hour. The data from each base station are used to calculate the average and standard deviation of packets per hour during each hour of the day. The maximum traffic volume in packets for any base communication package for each hour is also graphed.

The traffic for each base communication package was added during each hour to determine the total traffic volume through the network. During the heaviest hour for total inbound traffic, Hour 0, the front-end processor/cluster controller handled 28,995 packets. The heaviest hour for outbound traffic was Hour 20 when 18,511 packets were routed by the controller. The graph in Figure 4 shows inbound and outbound total traffic by hour.

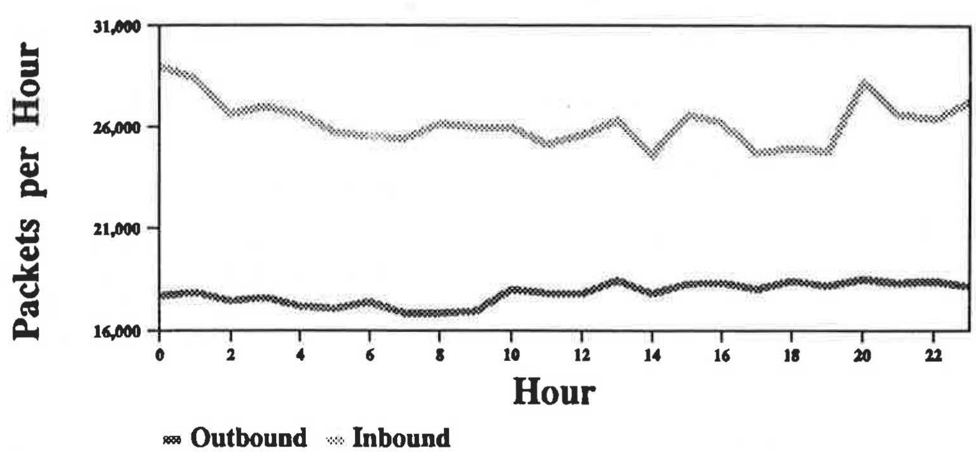

FIGURE 4 Total traffic.

Outbound traffic from each base station was a fairly constant 240 packets per hour. The maximum outbound traffic was in Hour 6 from Rocky Point, Oregon (near Portland), when 693 packets were transmitted. The second highest figure was for the same hour through North Platte, Nebraska, when 568 packets were transmitted. Figure 5 shows the average, standard deviation, and the maximum packets per hour outbound for each hour of the day.

The variation of inbound traffic was much greater than outbound. The busiest inbound base station was Trail Creek, Idaho (near Pocatello), during Hour 23 with 2,336 packets received. The second busiest was Black Mountain, Idaho (near Bonners Ferry), during Hour 20 when 2,238 packets were received. Figure 6 shows the average, standard deviation, and the maximum packets per hour inbound for each hour of the day.

The average outbound traffic volume of 240 packets per hour shown in Figure 5 is easily explained. ATCS specifications require each base communication package to transmit a cluster controller ID every 15 sec if there is no other traffic. This results in a baseline outbound load of 240 packets per hour.

The inbound traffic volume shown in Figure 6 varies greatly because locomotives move between base stations and base stations that provide coverage for yards have more locomotives within their coverage area. The inbound traffic is primarily from packets that do not directly support WORS or 3270 pass-through. ATCS specifications require each mobile communication package to transmit once every 5 min if it has no other traffic. This results in 12 packets per hour from every locomotive equipped with a mobile communication package. On-board terminals in WORS send a "WAMI" message every 2 min while in RF coverage, which results in 90 packets per hour from each locomotive with a terminal installed. These automatic messages are only sent if there is no other traffic during the time interval.

Traffic for a Job

The traffic to and from a "typical" WORS job is determined by monitoring randomly selected jobs. The AMCI front-end processor provides a function called TRACE, which monitors traffic to and from a selected locomotive. TRACE is started on a specific locomotive to record the traffic to and from that

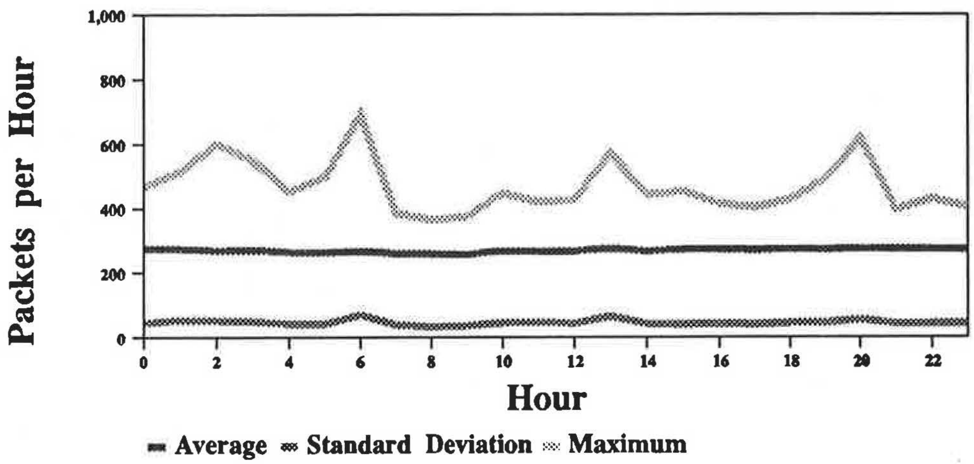

FIGURE 5 Outbound traffic per base communication package.

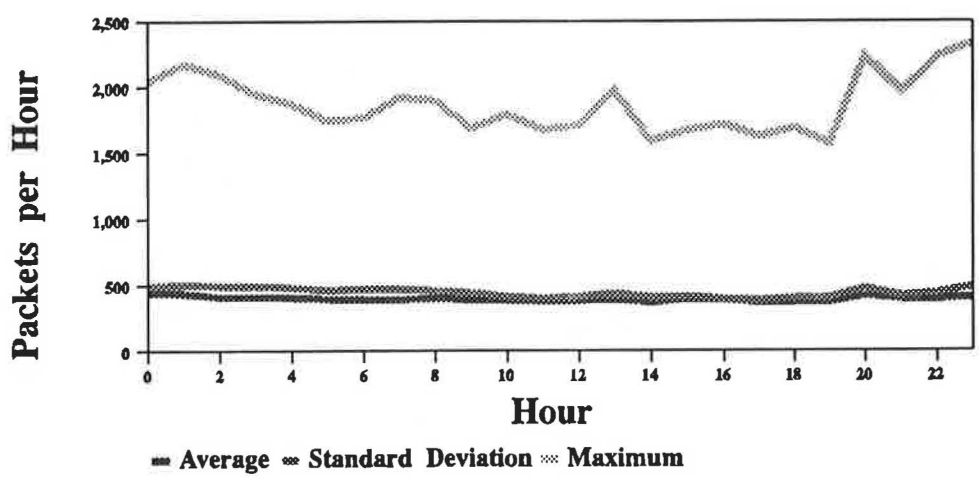

FIGURE 6 Inbound traffic per base communication package.

locomotive. The TRACE file is then converted to dBase for further analysis. This permits locomotive traffic to be processed and counted. When multiple packets are received at the front-end processor/cluster controller due to overlapping RF coverage, duplicate packets are removed when counting mobile packets per hour. Figure 7 summarizes the inbound traffic characteristics from 15 jobs. Figure 8 summarizes the outbound traffic characteristics from the same 15 jobs. The graphs show the minimum, maximum, and average packets per hour from the sampled jobs.

ANALYSIS

The average WORS traffic outbound for each base station uses less than 5 percent of the capacity of the base communication package. The busiest base used only 12 percent of the available outbound capacity. The heaviest load is on the inbound RF channel. The average inbound load is 500 packets per hour, conservatively 12 percent of the base communication packages' minimum inbound capacity of 4,200 packets per hour based on 256-byte packets. The maximum inbound traffic volume in an hour was 2,336 packets. Even during peak loads for WORS, almost half of the capacity of each base station is still available for operation of ATCS. Smaller-sized packets will result in a higher capacity in packets per hour. Because the actual traffic is a mixture of small and large packets, this analysis of capacity is conservative.

The equipment in an ATCS communication system automatically generates traffic to verify that it is in communication with the central system. Base stations send a packet every 15 sec. On-board terminals originate a message every 2 min. At the present time, this type of traffic represents the majority of the load on the network. When "real" messages are sent during a time interval, the automatic message is not sent. This automatically generated traffic establishes a minimum traffic level for the network. Initial application traffic merely replaces the automatically generated traffic and does not result

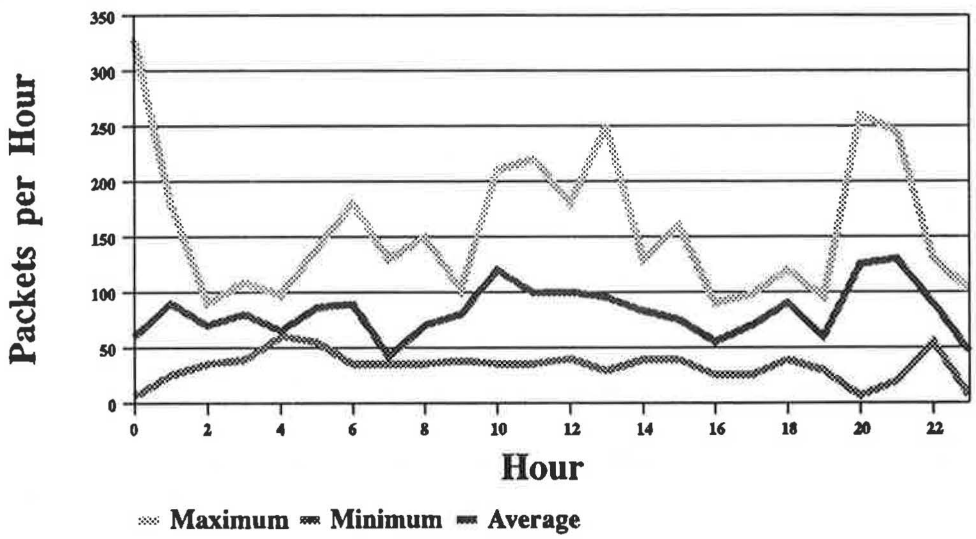

FIGURE 7 Inbound traffic per job.

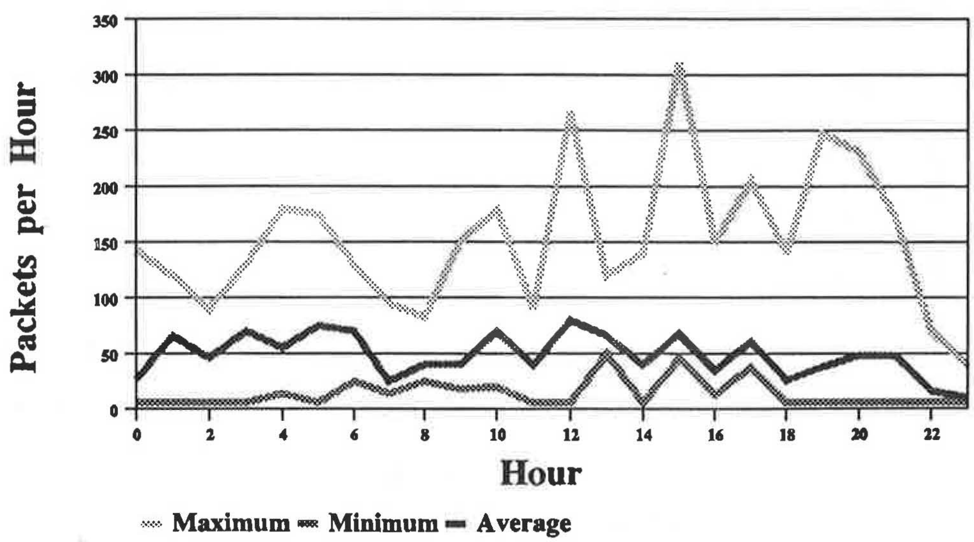

FIGURE 8 Outbound traffic per job.

in much increase in network traffic. Beyond this minimum traffic level, additional application traffic will result in increased loads on the network.

The successful large-scale implementation of Specification 200 on Union Pacific indicates that the ATCS specifications provide the basis for a viable communication network. A Specification 200 communication system can provide reliable communications for a variety of applications including ATCS, work order reporting, locomotive health monitoring, code line replacement, and maintenance-of-way management. The RF network can be engineered through site selection and use of multiple channels to provide the needed capacity for a railroad's data communication requirements.

The ATCS communication system is designed to support other applications in addition to train control. To ensure that critical traffic and low-priority traffic can share the same communication system, a priority value is assigned to each message. The priority determines the retry interval for the message and causes the message to pass lower-priority messages in any queue. The ATCS protocol suite establishes eight priority levels with two subpriority levels (RF ARQ mechanism disabled or enabled) per priority for a total of 16 "channel groups." The eight priority levels permit traffic to move through the network according to its relative importance. Packets for emergency messages have the highest priority. Packets involving train control have a higher priority than WORS packets or data base download packets.

A network handling multiple types of traffic will be engineered so that the delay requirements established in Section 3.2.1.1 of Specification 200 are met. According to these guidelines, the delay for 99 percent of emergency traffic between vehicle and wayside will be 4 sec or less, and dispatch system to a vehicle will be 10 sec or less. Normal traffic will have an average delay of 30 sec and a 99 percent delay of 225 sec. Operational traffic will have an average delay of 10 sec and a 99 percent delay of 75 sec.

SUMMARY

AMCI has been working with Union Pacific Railroad to implement a communication network based on ATCS Specification 200. This paper presents an estimate of the expected capacity (throughput) of the network based on analytical models. For convenience, throughput is estimated in packets per hour. For the largest packet size, 256 bytes, approximately 5,600 packets can be sent from a base station in 1 hr. Assuming a 256-byte packet, a lower limit of inbound capacity of a base station is estimated to be approximately 4,200 packets per hour.

The current loading level of the installed network is compared with the expected capacity. As many as 29 locomotives are operating through a given base communication package. The maximum outbound traffic from a base station was Rocky Point, Oregon, when 693 packets were transmitted during 1 hr. The busiest inbound base communication package was Trail Creek, Idaho, with 2,336 packets received in 1 hr. The average WORS traffic outbound for each base station uses less than 5 percent of the capacity. The average WORS inbound load uses, conservatively, 12 percent of the inbound capacity.

A Specification 200 communication system can provide reliable communications for a variety of applications including ATCS, work order reporting, locomotive health monitoring, code line replacement, maintenance-of-way management, and other future applications. To ensure that critical traffic and low-priority traffic can share the same communications system, a priority value is assigned to each message. The eight priority levels permit traffic to move through the network according to its relative importance. The RF network can be engineered through site selection and use of multiple channels to provide the needed capacity for a railroad's data communication requirements.

The successful large scale implementation of a Specification 200 network on Union Pacific indicates that the ATCS specifications provide the basis for a viable system capable of supporting ATCS, WORS, locomotive health, and other applications.

REFERENCES

1. Coll, Sheikh, Ayers, and Baily. The Communication System Architecture of the North American Advanced Train Control System. *IEEE Transactions on Vehicular Technology,* Vol. 39, No. 3, August 1990.
2. Sheikh, Coll, Ayers, and Baily. ATCS: Advanced Train Control System Radio Data Link Design Considerations. *IEEE Transactions on Vehicular Technology,* Vol. 39, No. 3, August 1990.
3. Sharif and Furman. Analytical Model for ATCS Inbound RF Channel Throughput. *Proceedings of the 1991 IEEE Vehicular Technology Conference.*
4. Kleinrock and Tobagi. Packet Switching in Radio Channels: Part I —Carrier Sense Multiple-Access Modes and Their Throughput Delay Characteristics. *IEEE Transactions on Communications,* Vol. COM-23, December 1975.
5. Mautschke, Furman, and Decker. Mobile Data Transmission in a Railroad Environment. *Proceedings of the 1990 IEEE Vehicular Technology Conference.*