Mobile-ip Aeronautical Network Simulation Study

William D. Ivancic and Diepchi T. Tran
Glenn Research Center, Cleveland, Ohio

National Aeronautics and
Space Administration

Glenn Research Center

June 2001

MOBILE-IP AERONAUTICAL NETWORK SIMULATION STUDY

William D. Ivancic and Diepchi T. Tran
National Aeronautics and Space Administration
Glenn Research Center
Cleveland, Ohio 44135

SUMMARY

NASA is interested in applying mobile Internet protocol (mobile-ip) technologies to its space and aeronautics programs. In particular, mobile-ip will play a major role in the Advanced Aeronautic Transportation Technology (AATT), the Weather Information Communication (WINCOMM), and the Small Aircraft Transportation System (SATS) aeronautics programs. This report presents the results of a simulation study of mobile-ip for an aeronautical network. The study was performed to determine the performance of the transmission control protocol (TCP) in a mobile-ip environment and to gain an understanding of how long delays, handoffs, and noisy channels affect mobile-ip performance.

INTRODUCTION

Mobile Internet protocol (mobile-ip) is a routing protocol that allows hosts (and networks) to seamlessly "roam" among various ip subnetworks, an essential feature in many wireless networks. Mobile-ip can also be useful in wireless networks where the mobile node's attachment point to the network is changing owing to varying conditions in the wireless medium, even if the mobile node is not physically moving. Mobile-ip can also be used in a wired network where the mobile node simply wishes to maintain its network identity, as the mobile node is always contacted through its home ip address.

Three basic elements in mobile-ip are the home agent, the foreign agent, and the mobile node (ref. 1):

1. The home agent (HA) is a router on a mobile node's home network. The HA tunnels datagrams for delivery to the mobile node when it is away from home and maintains the mobile node's current location information.
2. The foreign agent (FA) is a router on a mobile node's visited network that provides routing services to the mobile node while registered. The FA detunnels and delivers datagrams to the mobile node that were tunneled by the mobile node's HA. For datagrams sent by a registered mobile node, the FA may serve as a default router.
3. The mobile node (MN) is a host or router that changes its attachment point from one network or subnetwork to another. An MN may change its location without changing its ip address. It may continue to communicate with other Internet nodes at any location by using its (constant) ip address, if link-layer connectivity to an attachment point is available.

A mobile node is always identified by its home ip address, regardless of its current Internet attachment point. When the mobile node moves to another subnetwork, it will ask the FA to act as its agent in communicating with the HA. If the FA can accommodate this, it will provide the mobile node with a care-of-address. A tunnel will be set up between the FA and the HA whereby the HA forwards to the care-of-address all messages sent to the mobile node's home ip address.

NASA'S INTERESTS

NASA is interested in applying commercial-off-the-shelf (COTS) Internet technology to NASA missions to reduce costs while simultaneously upgrading its communications networks. Applying mobile-ip technologies to NASA missions will facilitate these goals.
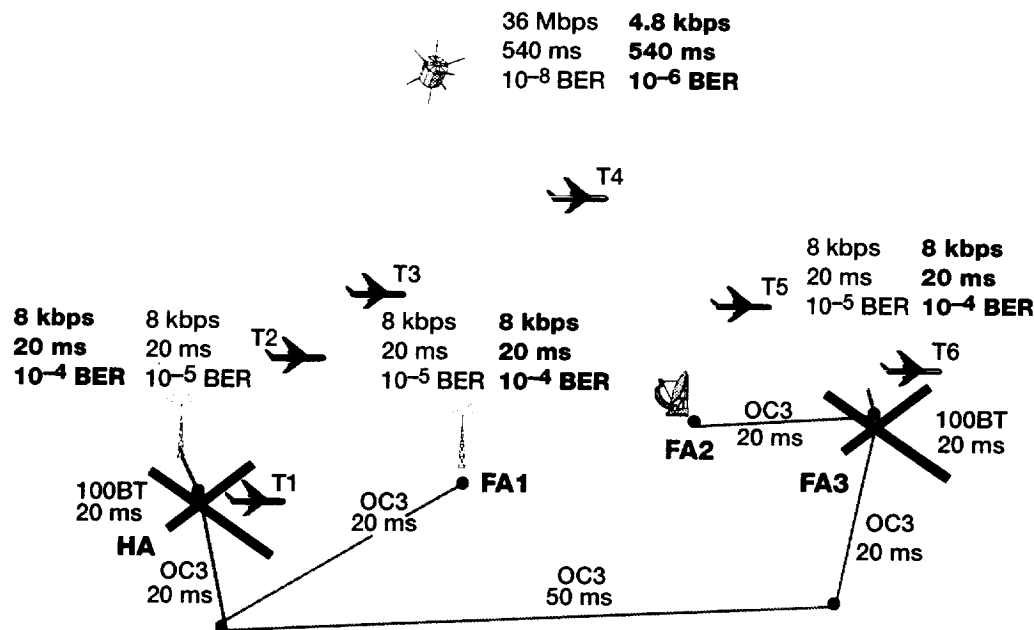
Figure 1.—Proposed simulated aeronautic network.

NASA has numerous applications where mobile-ip is desired: general computing; aeronautical telecommunications networks; orbiting space science missions; and terrestrial science missions (Earth, Moon, Mars, asteroids, etc.). Spacecraft communicating through several ground terminals networked together require mobile-ip (assuming duplex communications) because each ground terminal is an independent node on the network. An aeronautical network requires mobile-ip to maintain connectivity. As the aircraft moves from region to region, it traverses various subnetworks, one for each airport or air traffic control center (fig. 1). Thus, mobile-ip will play a major role in NASA's aeronautics programs including the Advanced Aeronautic Transportation Technology (AATT), the Weather Information Communication (WINCOMM), and the Smart Aircraft Transportation System (SATS) programs (refs. 2 to 4).

PURPOSE

We investigated four major questions in this study:

1. What are the mobile-ip registration and file transfer times?
2. How are handoffs handled and how does delay affect them?
3. How do errored links affect the mobile-ip protocol?
4. What parameters are critical to monitor in real-world mobile-ip networks?

This report addresses all four questions.

PROPOSED SIMULATION CONFIGURATION

Figure 1 shows the proposed aeronautical network configuration we simulated. We conceived that initially (time T1) the aircraft (mobile node) would be attached to its HA by some type of umbilical cord (wired network). Thus, the link would be error free and have a rate of 100 Mbps and a very low delay of 40-ms round-trip time. At time T2 the mobile node would no longer be connected by a wired network but would be attached to the HA by a very high-frequency data link (VDL). The VDL is a noisy, low-rate link. At time T3 the mobile node would be con-

nected to the first foreign agent (FA1) by a VDL link. At time T4 the mobile node would be connected to FA2 by a satellite link. Thus, there would now be a high delay in the link. Depending on the simulation setup the satellite link may be either near error free or noisy, and it may be either high bandwidth or low bandwidth. At time T5 the mobile node would be connected to FA3 by a VDL terminal. Finally, at time T6 the mobile node would be connected to FA3 by a wired network.

SIMULATION TOOLS

We chose Network Simulator (ns) software as our simulation tool (ref. 5). The ns was developed under the Defense Advanced Research Projects Agency (DARPA)–funded Virtual InterNetwork Testbed research project. The project's aim was to build a network simulator for the study of scale and protocol interaction in current and future network protocols. The ns is the de facto standard used to evaluate the TCP. Many extensions have been added to ns to accommodate mobile-ip. Because the ns source code was available, we could determine any short-comings to current protocol implementations and add to the overall research knowledge base and tool sets.

NS SIMULATION SCENARIO

We used the March 6, 2000, daily-snapshot version 2.1b6 of ns. This version did not have an error model that would work directly with a wireless link. Therefore, we set the bit error rate (BER) on wired links connected to the HA and the FA's before going to a wireless channel. In addition, the ns required a mobile node to be configured as either wired or wireless. We configured it as a wireless node. We started the simulation at time T2 and stopped at time T5.

Figure 2 shows the diagram that we used for the simulation. Our simulation consisted of one MN, one HA, and three FA's. The links are labeled with their bandwidth capacity, BER, and delay. A uniform distributed random number generator was used to inject the errors onto the links. For a nondeterministic error distribution, the seed was generated on the basis of the current time of day and a counter. All nodes marked as W1 to W6 in figure 2 were configured as wired nodes. The location of the HA, the FA's, and the initial MN position are indicated in x,y coordinates.
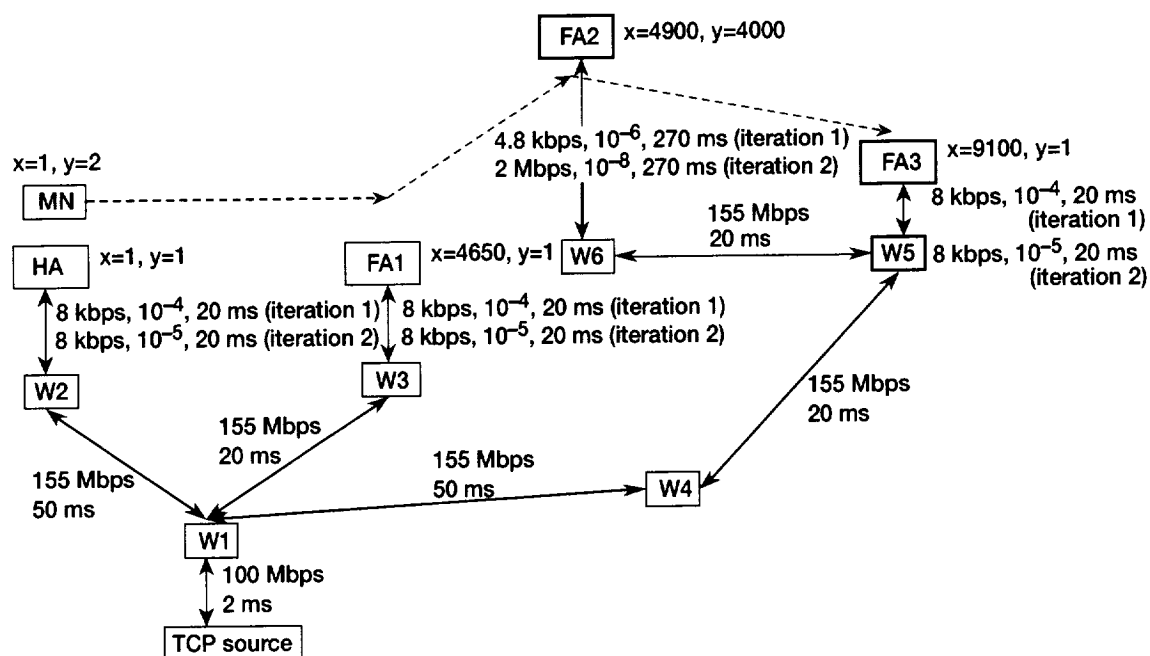

Figure 2.—ns simulation diagram.

We used the Sack TCP as indicated in RFC2018 (ref. 6) in our simulation. All wireless and TCP parameters were set to ns default values[1] except for the TCP packet size, the receiver's window, and the transmitted and received gain of the antenna. We set the TCP packet size to 512 bytes and the advertisement window to 5000 packets so that it was not a limiting factor in transfers.

The 36-Mbps satellite channel could not be utilized fully because IEEE802.11 link model used in the wireless domain was limited to 2-Mbps bandwidth, effectively limiting our satellite bandwidths to 2 Mbps. The transmitted and received antenna gains were set to 40.0, making the nominal range of the MN, HA, and FA's 1581 m. The destination-sequence distance-vector (DSDV) routing protocol (ref. 7) was used to route packets between the MN, HA, and FA's. Fifty seconds into the simulation the MN started to move away from the HA at 200 m/s. The MN paused for 20 s in the range of FA1 and for 40 s in the range of FA2 before it stopped in the range of FA3. The TCP source also started to send 512-byte packets by file transfer protocol (FTP) to the MN 50 s into the simulation (as soon as the MN started moving).

ASSUMPTIONS

For wired links we assumed a BER of zero and no congestion. For the VDL terminal connections we assumed an 8-kbps data rate and error rates of either 10^{-4} or 10^{-5}, typical for VDL channels. For the satellite connections we assumed either a near-error-free, high-rate channel (2 Mbps and 10^{-8} BER) or a noisy, low-rate channel (4.8 kbps and 10^{-6} BER), typical of today's satellite links.

VARIABLES

To determine network performance, we performed TCP transfers while transitioning the network. Because TCP performance is highly dependent on when a packet is lost during a session, we needed to perform multiple runs with a variety of file sizes (refs. 8 and 9). The files sizes we chose were 10 kbytes, 50 kbytes, and 1 Mbyte. TCP performance is also highly dependent on BER. Thus, we needed to investigate performance over various BER links as described in the assumptions. In addition, TCP performance can be affected by the packet size owing to packet-size interactions with BER and with the buffer queue sizes within the network. Thus, we initially ran simulations for the following packet sizes: 512, 1000, and 1500 bytes. The 1000 bytes is the ns default and 1500 bytes is typical of an Ethernet packet. However, when we used packets of 1000 bytes or greater, we noted some strange behavior at the HA as described in the section Possible Bugs in Current Software Packet. We suspect this behavior resulted from a bug in the mobile-ip implementation code. This anomaly rendered simulation with large packet sizes (greater than 1000 bytes) questionable. Thus, we ran all further simulations using only a 512-byte packet size. Table I summarizes the variables used in the simulation. We ran 20 iterations; the first iteration was for high-BER, low-bandwidth. The second iteration was for low-BER, high-bandwidth.

The following sections describe the behavior of the MN, HA, and FA's, the results of the average throughput from 30 test runs for each file size, and possible bugs in current software packages and identify areas that need to be further developed.

TABLE I.—SIMULATION VARIABLES

File size, bytes	Iteration	Channel	BER	Bandwidth
10k, 50k, 1M	1st	VDL	10^{-4}	8 kbps
		Sat	10^{-6}	4.8 kbps
10k, 50k, 1M	2nd	VDL	10^{-5}	8 kbps
		Sat	10^{-8}	2 Mbps (due to 802.11 implementation limitation)

---

[1]The TCP default parameters are explained in chapter 28, section 28.1.4 of reference 5. Default values used for the wireless simulation described herein are listed in appendix A.

OBSERVATIONS

In the trace file obtained from a simulation, the ns did not use the Internet control message protocol (ICMP) and the user data protocol (UDP) to distinguish the advertisement and registration messages. Rather it used the UDP for both of these messages. Therefore, from the definition of advertisement and registration in reference 10 (chapters 3 and 4), the broadcast messages are advertisement messages and the others are registration. In the ns the HA, FA's, and MN send out the broadcast message to advertise their media access control (MAC) addresses every 0.5 s. When the MN moves into an FA's range, it also sends a registration message to that FA every 0.5 s or every time the FA's MAC address is received.

In our simulation, besides the dropping of packets caused by error and in the handoff periods, some UDP packets in the trace file were dropped with address resolution protocol (ARP) and drop_RTR_MAC_call back (CBK) flags. Current ARP is implemented in the ns so that when a packet sent to an MN arrives at an HA or an FA before the MN's MAC address is learned, it is buffered. If an HA or an FA receives a new packet also sent to the MN and the MAC address of the MN still is not resolved, the buffered packet is dropped. Therefore, in the trace file we saw some droppings of UDP packets with ARP flags.

When the MN was already in the 1581-m range, some droppings still occurred because base stations (HA and FA's) could not locate the MN at that moment. This type of packet dropping is indicated by a CBK flag in the trace file. These droppings happened because the DSDV routing protocol was used. In DSDV routing the routing packets are exchanged between neighboring nodes (MN, HA, and FA), and the routing updates may be triggered or routine. The update in an FA routing table could be triggered when the MN moved into that FA's range. If the FA routing table was not updated every three minimum update periods, the MN was declared unreachable and the packet was dropped. Therefore, in our trace file while waiting for routing table update, the FA's sometimes dropped the third and fourth UDP packets with CBK flags after dropping every two UDP packets with ARP flags. With all BER's set to zero, it took 16 to 26 s from the time the MN learned an FA's MAC address to the time it received the first TCP packet through that FA.

Figure 3 shows packet transfer through the satellite channel with error-free links in the first and second iterations (ref. Table 1). Although the IEEE802.11 bandwidth is limited, the packets still ramped up better in the second iteration with satellite channel bandwidth (effectively 2 Mbps) than in the first iteration with a 4.8-kbps bottleneck channel. Only 30 packets (44th to 74th) were received and acknowledged (ACKed) through the 4.8-kbps link in 38 s (from 122 to 160 s) (fig. 3(a)). In the same time interval 57 packets were received and ACKed through the satellite link (fig. 3(b)).

SIMULATION RESULTS

The results presented in this report are valid only for our simulation. They will differ depending on the simulation scenario. Figure 4 shows the average throughput of 30 test runs for each file size. As we expected, in the same channel condition the average throughput of the bigger file size was better than the average throughput of the two smaller file sizes. In our simulation one reason for this difference was that the smaller files had a larger percentage of handoff delay periods over the total transmission time than the bigger file. A handoff delay period is the time when an MN leaves the domain of an HA or an old FA to the time when it receives the first TCP packet from a new FA. The smaller files also spent more time in the slow-start phase, where the average congestion window is small. The bigger file spent more transmission time in the congestion-avoidance phase, where the average congestion window is big. Without the BER effect, as shown in figure 5, the times that the MN spent in handoff delay periods in the 10-kbyte, 50-kbyte, and 1-Mbyte files were 63, 52, and 7 percent, respectively, of the total transmission time.

As shown in figure 4, the average throughput of 1-Mbyte and 10-kbyte file transfers was about 2.4 times better in the second iteration than in the first. The average throughput of 50-kbyte file transfers was 3.4 times better in the second iteration than in the first. Obviously, because of the lower BER of all link, the average throughputs of all file sizes were higher in the second iteration than in the first. The 50-kbyte file showed significant improvement in the second iteration versus the first because 21 of 30 transfers were finished before the MN moved to FA3. For the 1-Mbyte file the MN received the last packet through FA3 in both iterations. For the 10-kbyte file all transfers in the second iteration and about 20 in the first iteration were finished when the MN was at FA1.
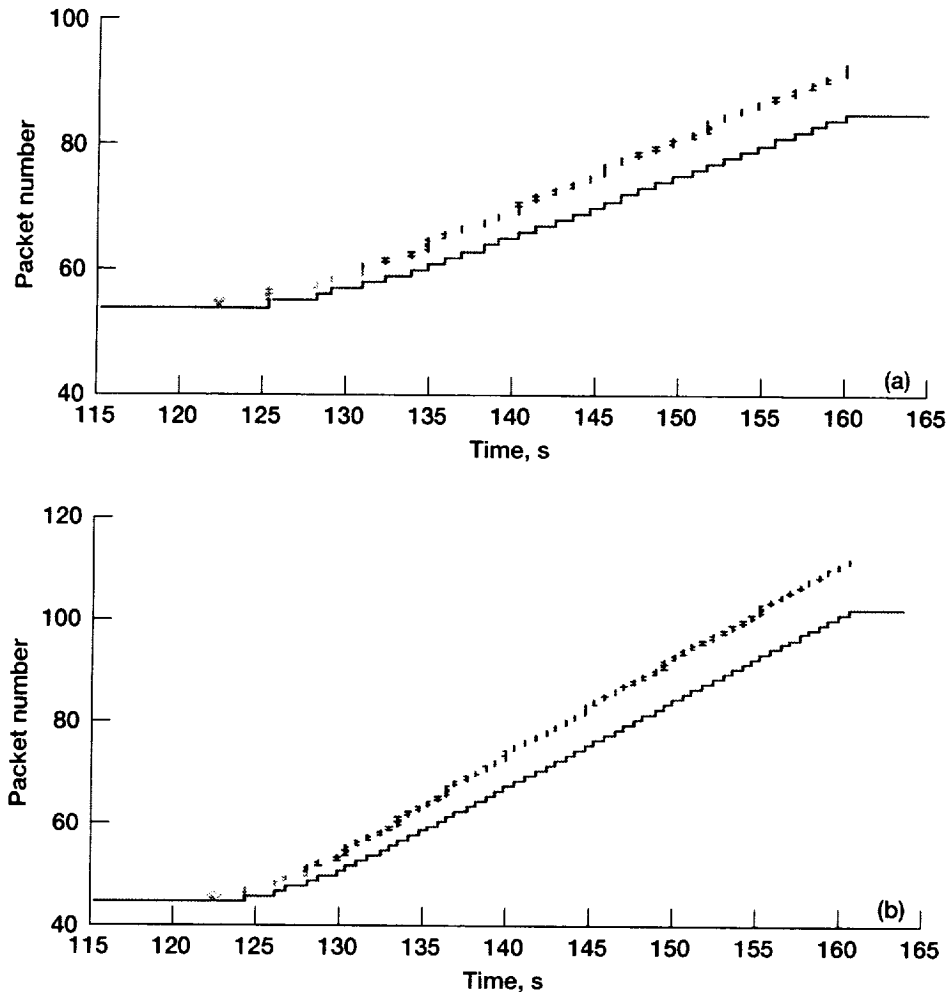
Figure 3.—Packet transfer. (a) Through 4.8-kbps channel. (b) Through 2-Mbps channel.
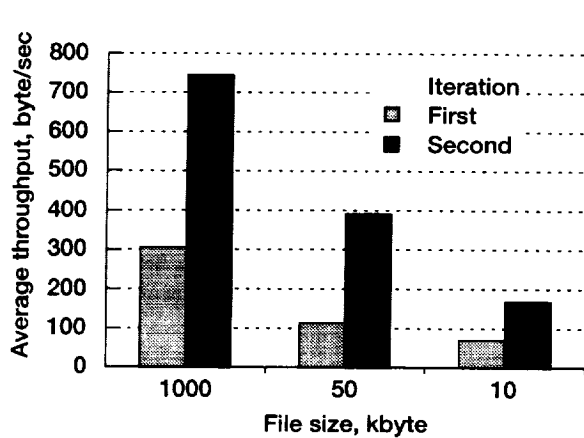

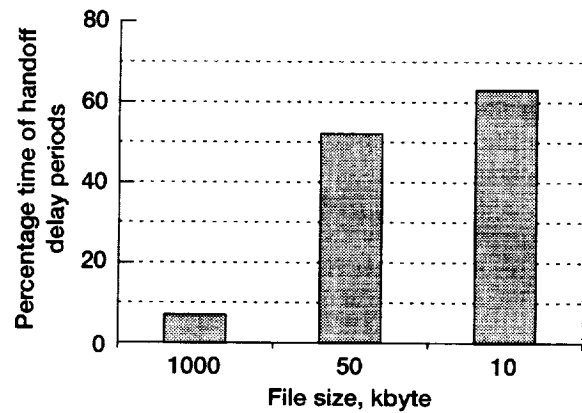Figure 4.—Average throughput as function of file size.

Figure 5.—Percentage of total transmission time in handoff delays as function of file size.

To investigate the effect of handoff delays on throughput, we ran one test for each file size in both iterations of channel condition without errors inserted in any links. Total handoff delays took 23 s for the 10-kbyte file and 85.7 s for the 50-kbyte and 1-Mbyte files in the first iteration and 23, 46, and 79 s, respectively, in the second iteration. Figure 6 and 7 show a trace of a 50-kbyte file transfer of the first iteration run error-free and with with errors. Figure 6 shows the retransmission timers expiring during handoffs, and the TCP entered the slow-start phase after them.

Figure 7 shows the handoffs took 25 s between the HA and FA1, 77.2 s between FA1 and FA2, and 100 s between FA2 and FA3, for a total delay of 202.2 s. Figure 7 also shows that the delays caused by BER after the handoffs were small relative to the handoff delays.

In our simulation the handoff delay was a total of three delays: out-of-sight delay, mobile-ip (Mip) delay, and TCP delay. The out-of-sight delay happened when the MN moved out from an old base-station (HA or FA's) domain but was not yet in a new base-station domain. In the Mip delay, latency was caused by advertisement, solicitation, and registration procedures between base stations and the MN. TCP delay was the time that the TCP sender waited for the retransmission timer's times out before retransmitting its packets, even though the connection between the MN and the new base station was already available. In our six error-free simulations (the runs of three file sizes in two iterations), the out-of-sight delay took 25 to 32 percent of the total handoff delay, the Mip delay 40 to 60 percent, and the TCP delay 2.5 to 15 percent. Therefore, the Mip and out-of-sight delays were the first and second factors to degrade system performance efficiency.
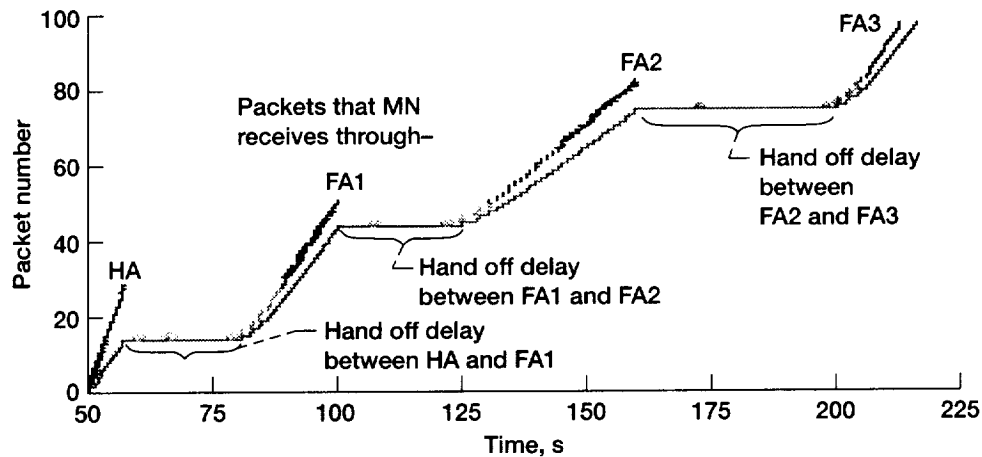

Figure 6.—50-kbyte file transfer in first iteration with error-free links.
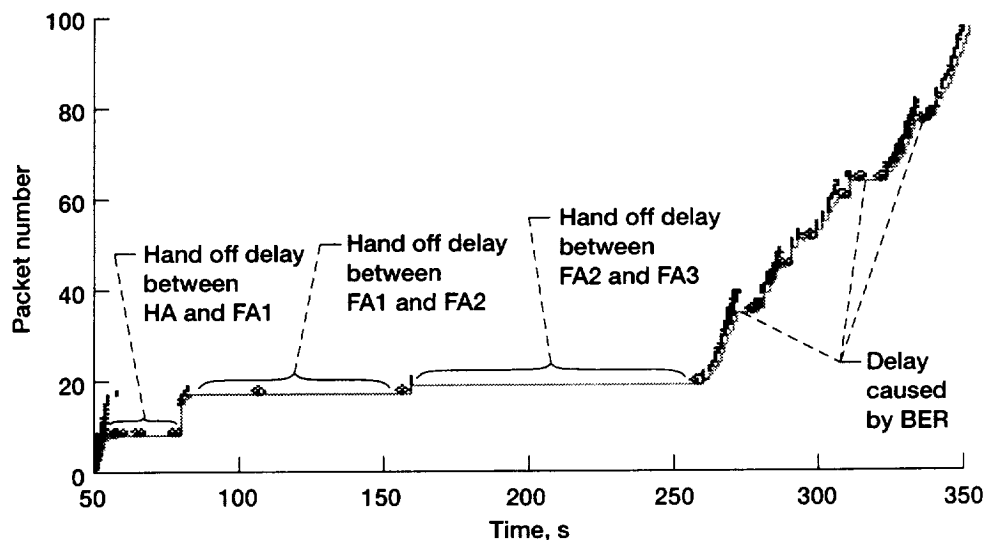

Figure 7.—50-kbyte file transfer in first iteration with errors.

A make-before-break mechanism, which allows an MN to complete the registration procedure with a new FA before being disconnected from an old FA, can be used to reduce Mip delays. However, this mechanism can be applied only when the MN can see both the old and new FA at the same time. The out-of-sight delay can be eliminated if base-station domains can fully cover the MN's moving path.

POSSIBLE BUGS IN CURRENT SOFTWARE PACKET

While working on the simulation we faced a problem when using 1000- and 1500-byte TCP packets. After an MN moved away from an HA into an FA's range, a TCP packet was sent to the routing level (RTR level) while the packets before and after it were encapsulated and forwarded to the FA by the HA. To single out the problem, we ran an error-free simulation with channel bandwidth set up as in the first iteration. Figure 8 shows part of the 1-Mbyte file transfer when an MN moved into FA3's range. The time-sequence plot shows that the 20th and 21st packets were sent to the MN at the same time. However, only the 20th packet was encapsulated, forwarded to FA3, and received by the MN. When the 21st packet arrived at the HA, it was sent up to the HA's RTR level. Then the acknowledgment (ACK) of the 20th packet triggered the TCP source to send the 22nd packet. The HA received, encapsulated, and forwarded the 22nd packet to FA3. Appendix B is a part of the trace file that shows this HA behavior.

As shown in figure 8, the HA behaved the same when it received packets 24, 27, 30, 33, etc., as it did with the 21st packet. We also observed this HA behavior when using the 1500-byte TCP packet. We think it may be caused by a bug in the ns code.

We also encountered a problem when using a default random number generator within ErrorModel. The default generator does not actually provide random errors. Therefore, we created the random errors by using the random number generator described in chapter 20 of reference 5 .

FUTURE WORK

The ns needs to be modified to create a simulation closer to a real network. Currently, no error model is associated with a wireless link. To have a BER in a wireless channel, we need to add a time-varying error model between the MAC and the link layer of a wireless node (MN) or base-station nodes (HA and FA's). In addition, the ns implementation used in our simulations did not support a smooth-handoff option (ref. 11). To study the effect of smooth-handoffs on the throughput, we will invistigate an extension to ns mobile-ip implemented by Joerg Widmer (ref. 12).

In order to more accurately model the 36-Mbps satellite channel, a satellite data link-layer model for the satellite link needs to be deployed rather than the IEEE802.11 data link.
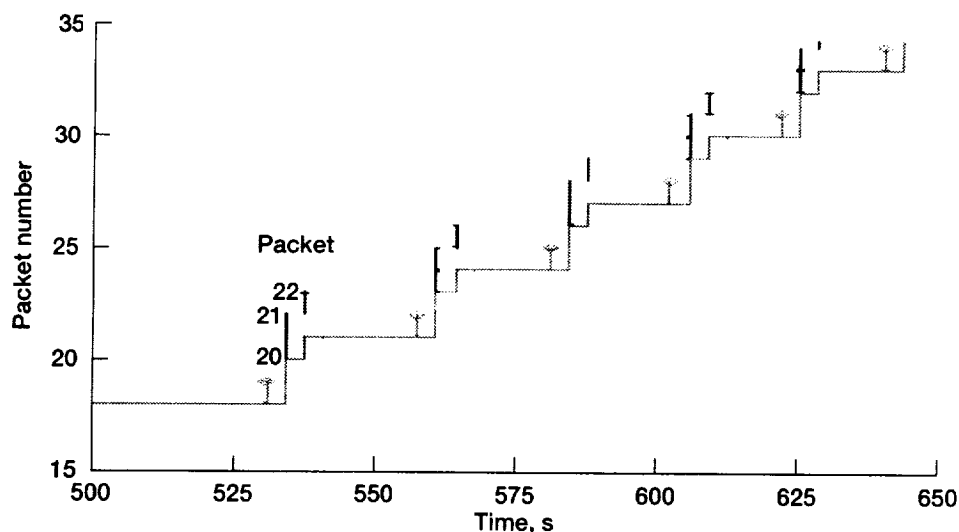

Figure 8.—1-Mbyte file transfer with 1000-byte packet size.

CONCLUSIONS

We have studied the effect of handoff delays and bit error rate (BER) on the throughput in a network consisting of wireless and wired domains. Our simulation consisted of file transfer protocol (FTP) over transmission control protocol (TCP) with mobile Internet protocol (mobile-ip). In our simulation we used three file sizes and two BER iterations and the bandwidth conditions for very high-frequency data link (VDL) and satellite channels. Without the BER effect our results showed that handoff delays took 63 and 52 percent of total transmission time in 10- and 50-kbyte files, respectively. When errors were inserted in the links, these handoff delays took even longer because advertisement, registration, and solicitation packets were dropped. The delays caused by BER were small relative to the handoff delays. Therefore, handoff delay has a critical impact on the throughput, especially for small file transfers. Besides the BER effect, handoff delay also depended on how the moving path of a mobile node was covered (fully or partially) by base-station domains and routing protocol. Using the make-before-break mechanism should greatly improve throughput efficiency.

REFERENCES

1. Perkins, C.E.: IP Mobility Support. Request for Comments 2002, Network Working Group, Oct. 1996.
2. Advanced Air Transportation Technology Program: *http://www.asc.nasa.gov/aatt/*, accessed Jan. 31, 2001.
3. Weather Information Communication Program: *http://wincomm.grc.nasa.gov/about/*, accessed Jan. 31, 2001.
4. Small Aircraft Transportation System Program: *http://www.unomaha.edu/~unoai/sats/*, accessed Jan. 31, 2001.
5. The Network Simulator: *http://www.isi.edu/nsnam/ns/* or *http://ee.lbl.gov/ns/*, accessed Jan. 31, 2001
6. Mathis, M.; et al.: TCP Selective Acknowledgment Options. Request for Comments 2018, Network Working Group, Oct. 1996.
7. Perkins, C.E.; and Bhagwat, P.: Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers. Comput. Commun. Rev., vol. 24, no. 4, 1994.
8. Glover, D.R.; and Kruse, H.: TCP Performance in a Geostationary Satellite Environment. Sect. 6, Wireless. Annual Review of Communications 1998, International Engineering Consortium, 1998.
9. Partridge, C.: Study and Simulation of Enhancements for TCP Performance Over Noisy High Latency Links. NASA/CR—1999-209167, 1999. (Also, *http://gltrs.grc.nasa.gov/GLTRS*, accessed Jan. 31, 2001.)
10. Perkins, C.E.: Mobile IP Design Principles and Practices. Addison-Wesley, Reading, MA, 1998.
11. Perkins, C.E.; and Johnson, D.B.: Route Optimization in Mobile IP. Network Working Group, Feb. 25, 1999.
12. Widmer J.: Extensions to the ns Network Simulator. *http://www.icsi.berkeley.edu/~widmer/mnav/ns-extension/*, accessed Jan. 31, 2001.

APPENDIX A
DEFAULT PARAMETERS

The default parameters used for the wireless simulation described herein are listed here and explained at *http://www.monarch.cs.cmu.edu.*

Simulator set AgentTrace_ ON
Simulator set RouterTrace_ OFF
Simulator set MacTrace_ OFF

LL set mindelay_ 50us
LL set delay_ 25us

Queue/Droptail/PriQueue set Prefer_Routing_Protocols 1

Antenna/OmniAntenna set X_ 0
Antenna/OmniAntenna set Y_ 0
Antenna/OmniAntenna set Z_ 1.5
Antenna/OmniAntenna set Gt_1.0
Antenna/OmniAntenna set Gr_1.0
Phy/WirelessPhy set CPThresh_ 10.0
Phy/WirelessPhy set CSThresh_ 1.559e-11
Phy/WirelessPhy set RXThresh_ 3.652e-10
Phy/WirelessPhy set Rb_ 2*1e6
Phy/WirelessPhy set Pt_ 0.2818
Phy/WirelessPhy set freq_ 914e+6
Phy/WirelessPhy set L_ 1.0
Phy/WirelessPhy set bandwidth_ 10e6

/* TCP source sends out 20th and 21th packets to HA */

```
+ 534.555449 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
- 534.555449 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
+ 534.555449 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
- 534.555529 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
 r 534.557529 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
+ 534.557529 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
- 534.557529 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
 r 534.557609 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
+ 534.557609 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
- 534.557609 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
 r 534.57758 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
+ 534.57758 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
- 534.57758 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
 r 534.57766 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
+ 534.57766 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
- 535.57758 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
```

/* HA encapsulates and forwards the 20th packet. But it sends the 21th packet to the RTR level */

```
 r 535.59758 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 20 3231
+ 535.59758 7 2 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 20 3231
- 535.59758 7 2 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 20 3231
 r 536.59758 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 21 3232
 r 536.597580443 _7_ RTR — 3232 tcp 1000 [0 0 0 0] ———— [0:0 20971521:2 290] [21 0] 0 0
 r 536.63758 7 2 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 20 3231
+ 536.63758 2 1 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 20 3231
```

/* TCP source sends out the 22nd packet after it receives ACK of the 20th packet */

```
- 537.923869 1 0 ack 60 ———— 1 5.0.1.2 0.0.0.0 20 3254
 r 537.925873 1 0 ack 60 ———— 1 5.0.1.2 0.0.0.0 20 3254
+ 537.925873 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
- 537.925873 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
 r 537.927953 0 1 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
+ 537.927953 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
- 537.927953 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
 r 537.948005 1 2 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
```

/* HA encapsulates and forwards the 22nd packet */

```
+ 537.948005 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
- 537.948005 2 7 tcp 1000 ———— 1 0.0.0.0 5.0.1.2 22 3256
+ 538.968005 7 2 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256
- 538.968005 7 2 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256
 r 540.008005 7 2 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256
+ 540.008005 2 1 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256
- 540.008005 2 1 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256
```

r 540.028058 2 1 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256
+ 540.028058 1 4 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256
- 540.028058 1 4 tcp 1020 ———— 1 0.0.0.1 8.0.0.1 22 3256

The trace file format is explained in chapter 21 of reference 5. In our trace file the node identifications of the TCP source, wired node W1, wired node W2, wired node W4, and the HA node are indicated by 0, 1, 2, 4, and 7, respectively. The ip address of the TCP source and the MN are 0.0.0 and 5.0.1, respectively.

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 2001	Technical Memorandum

4. TITLE AND SUBTITLE

Mobile-ip Aeronautical Network Simulation Study

6. AUTHOR(S)

William D. Ivancic and Diepchi T. Tran

5. FUNDING NUMBERS

WU-258-90-00-00

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
John H. Glenn Research Center at Lewis Field
Cleveland, Ohio 44135-3191

8. PERFORMING ORGANIZATION REPORT NUMBER

E-12679

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

NASA TM—2001-210751

11. SUPPLEMENTARY NOTES

Responsible person, Diepchi T. Tran, organization code 5610, 216-433-8861.

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified - Unlimited
Subject Category: 01   Distribution: Nonstandard

Available electronically at http://gltrs.grc.nasa.gov/GLTRS

This publication is available from the NASA Center for AeroSpace Information, 301-621-0390.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

NASA is interested in applying mobile Internet protocol (mobile-ip) technologies to its space and aeronautics programs. In particular, mobile-ip will play a major role in the Advanced Aeronautic Transportation Technology (AATT), the Weather Information Communication (WINCOMM), and the Small Aircraft Transportation System (SATS) aeronautics programs. This report presents the results of a simulation study of mobile-ip for an aeronautical network. The study was performed to determine the performance of the transmission control protocol (TCP) in a mobile-ip environment and to gain an understanding of how long delays, handoffs, and noisy channels affect mobile-ip performance.

14. SUBJECT TERMS

Satellite communications; Protocols; TCP; Mobile-ip

15. NUMBER OF PAGES

18

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102