

Tutorial Examples of Interprocess Communication in Berkeley UNIX 4.2 BSD

Stuart Sechrest

Computer Science Research Group
Computer Science Division
Department of Electrical Engineering and Computer Science
University of California, Berkeley

ABSTRACT

Berkeley UNIX 4.2BSD offers several choices for interprocess communication. To aid the programmer in developing programs comprising cooperating processes, the different choices are discussed and a series of example programs are presented. These programs demonstrate in a simple way the use of pipes, socketpairs, and sockets and the use of datagram and stream communication.

1. Goals

At times it is necessary to structure a program as a set of cooperating processes—perhaps because parts of the program work on separate machines. At times it is desirable to structure a program in this way—perhaps because the program has several distinct parts that can operate concurrently. Once it has been decided to implement a program as a set of processes, it is necessary to choose a method of communication. Berkeley UNIX 4.2BSD offers several choices. This paper presents simple examples that illustrate some of the choices. The reader is presumed to be familiar with the C programming language [Kernighan & Ritchie 1978], but not necessarily with the system calls of the UNIX system or with processes and interprocess communication. The paper reviews the notion of a process and the types of communication that are supported by Berkeley UNIX 4.2BSD. A series of examples are presented that create processes that communicate with one another. The programs show different ways of establishing channels of communication. Finally, the calls that actually transfer data are reviewed. To clearly present

UNIX is a trademark of AT&T Bell Laboratories.

This work was sponsored by the Defense Advanced Research Projects Agency (DoD), ARPA Order No. 4031, monitored by the Naval Electronics Systems Command under contract No. N00039-C-0235. The views and conclusions contained in this document are those of the author and should not be interpreted as representing official policies, either expressed or implied, of the Defense Research Projects Agency or of the US Government.

how communication can take place, the example programs have been cleared of anything that might be construed as useful work. They can, therefore, serve as models for the programmer trying to construct programs comprising cooperating processes.

2. Processes

A *program* is both a sequence of statements and a rough way of referring to the computation that occurs when the compiled statements are run. A *process* can be thought of as a single line of control in a program. Most programs execute some statements, go through a few loops, branch in various directions and then end. They are single process programs. Programs can also have a point where control splits into two independent lines, an action called *forking*. In UNIX these lines can never join again. A call to the system routine *fork()* causes a process to split in this way. The result of this call is that two independent processes will be running, executing exactly the same code. Memory values will be the same for all values set before the fork, but, subsequently, each version will be able to change only the value of its own copy of each variable. The only difference between the two will be the value returned by *fork()*. The parent will receive a process id for the child, the child will receive a zero. Calls to *fork()*, therefore, typically precede, or are included in, an if-statement.

A process views the rest of the system through a private table of descriptors. The descriptors can represent open files or sockets (sockets are communication objects that will be discussed below). Descriptors are referred to by their index numbers in the table. The first three descriptors have been given special names, *stdin*, *stdout* and *stderr*. These are the standard input, output and error. When a process forks, its descriptor table is copied to the child. Thus, if the parent's standard input is being taken from a terminal (devices are also treated as files in UNIX), the child's input will be taken from the same terminal. Whoever reads first will get the input. If, before forking, the parent changes its standard input so that it is reading from a new file, the child will take its input from the new file. It is also possible to take input from a socket, rather than from a file.

3. Pipes

Most users of UNIX know that they can pipe the output of a program "prog1" to the input of another, "prog2," by typing the command "*prog1 | prog2.*" This is called "piping" the output of one program to another because the mechanism used to transfer the output is called a pipe. When the user types a command, the command is read by the shell, which decides how to execute it. If the command is simple, say just "*prog1,*" the shell forks a process, which executes the program prog1 and then dies. The shell waits for this termination and then prompts for the next command. If the command is a compound command, say "*prog1 | prog2,*" the shell creates two processes connected by a pipe. One process runs the program prog1, the other runs prog2. The pipe is an I/O mechanism with two ends, or sockets. Data that is written into one socket can be read from the other.

Since a program specifies its input and output only by the descriptor table indices, which appear as variables or constants, the input source and output destination can be changed without changing the text of the program. It is in this way that the shell is able to set up pipes. Before executing prog1, the process can close whatever is at *stdout* and replace it with one end of a pipe. Similarly, the process that will execute prog2 can substitute the opposite end of the pipe for *stdin*.

Let us now examine a program that creates a pipe for communication between its child and itself (Figure 1). A pipe is created by a parent process, which then forks. When a process forks, the parent's descriptor table is copied into the child's.

In Figure 1, the parent process makes a call to the system routine *pipe()*. This routine creates a pipe and places descriptors for the sockets for the two ends of the pipe in the process's descriptor table. *Pipe()* is passed an array into which it places the index numbers of the sockets it created. The two ends are not equivalent. The socket whose index is returned in the low word of the array is opened for reading only, while the socket in the high end is opened only for writing. This corresponds to the fact that the standard input is the first descriptor of a process's descriptor table and the standard output the second. After creating the pipe, the parent creates the child

```

#include <stdio.h>

#define DATA "Bright star, would I were steadfast as thou art . . ."

/* This program creates a pipe, then forks.  The child communicates to the
 * parent over the pipe.
 * Notice that a pipe is a one-way communications device.  I can write to
 * the output socket (sockets[1], the second socket of the array returned
 * by pipe()) and read from the input socket (sockets[0]), but not vice versa.
 */

main()
{
    int sockets[2], child;

    /* Create a pipe */
    if ( pipe(sockets) < 0 )
        perror("opening stream socket pair");

    if (child = fork()) {
        char buf[1024];

        /* This is still the parent.  It reads the child's message. */
        close(sockets[1]);
        if ( read(sockets[0], buf, 1024) < 0 )
            perror("reading message");
        printf("-->%s\n", buf);
        close(sockets[0]);
    }
    else {
        /* This is the child.  It writes a message to its parent. */
        close(sockets[0]);
        if ( write(sockets[1], DATA, sizeof(DATA)) < 0 )
            perror("writing message");
        close(sockets[1]);
    }
}

```

Figure 1 Use of a pipe

with which it will share the pipe by calling *fork()*. Figure 2 illustrates the effect of a fork. The parent process's descriptor table points to both ends of the pipe. After the fork, both parent's and child's descriptor table point to the pipe. The child can then use the pipe to send a message to the parent.

Just what is a pipe? It is a one-way communication mechanism, with one end made for reading and the other end for writing. Therefore, parent and child need to agree on which way to turn the pipe, from parent to child or the other way around. Using a pipe for communication




Figure 2 Sharing a pipe between parent and child

from parent to child and from child to parent would be possible (since both processes have references to both ends), but very complicated. If the parent and child are to have a two-way conversation, the parent creates two pipes, one for use in each direction. (In accordance with their plans both parent and child in the example above close the socket that they will not use. It is not required that unused descriptors be closed, but it is good practice.) A pipe is also a *stream* communication mechanism, that is all messages sent through the pipe are placed in order and

delivered. When the reader asks for some number of bytes of this stream, he is given as many bytes as are available, up to the amount of the request. Note that this pays no attention to whether those bytes all came from the same call to *write()* or came from several calls to *write()* but were concatenated.

4. Socketpairs

Berkeley UNIX 4.2BSD provides a slight generalization of pipes. A pipe is a pair of connected sockets for one-way stream communication. One can also get a pair of connected sockets for two-way stream communication by calling the routine *socketpair()*. The program in Figure 3 calls *socketpair()* to create such a connection. The program uses the link for communication in both directions. Since socketpairs are an extension of pipes, their use resembles that of pipes. Figure 4 illustrates the result of a fork following a call to *socketpair()*. Socketpairs have only been implemented for stream communication in the UNIX domain. A domain is a space of names that may be bound to sockets and implies certain other conventions. The UNIX domain uses UNIX path names for sockets. It only allows communication between sockets on the same machine. *Socketpair()* takes as arguments a specification of a domain, a style of communication, and a protocol. These are the parameters shown in the example. Domains and protocols will be discussed in the next section. Note that the header files `<sys/socket.h>` and `<sys/types.h>` are required in this program. The constants `AF_UNIX` and `SOCK_STREAM` are defined in `<sys/socket.h>`, which in turn requires the file `<sys/types.h>` for some of its definitions.

5. Domains and Protocols

Pipes and socketpairs are a simple solution for communicating between a parent and child or between child processes. What if we wanted to have processes that have no common ancestor to communicate? Neither standard UNIX pipes nor socketpairs are the answer here, since both

```

#include <sys/types.h>
#include <sys/socket.h>
#include <stdio.h>

#define DATA1 "In Xanadu, did Kublai Khan . . ."
#define DATA2 "A stately pleasure dome decree . . ."

/* This program creates a pair of connected sockets then forks and
 * communicates over them. This is very similar to communication
 * with pipes, however, socketpairs are two-way communications objects.
 * Therefore I can send messages in both directions.
 */

main()
{
    int sockets[2], child;

    if ( socketpair (AF_UNIX,SOCK_STREAM,0,sockets) < 0 )
        perror("opening stream socket pair");
    if (child = fork()) {
        char buf[1024];
        close(sockets[0]);
        if ( read( sockets[1], buf, 1024,0) < 0 )
            perror("reading stream message");
        printf("-->%s\n", buf);
        if ( write(sockets[1],DATA2,sizeof(DATA2)) < 0 )
            perror("writing stream message");
        close(sockets[1]);
    }
    else {
        char buf[1024];
        close(sockets[1]);
        if ( write(sockets[0],DATA1,sizeof(DATA1)) < 0 )
            perror("writing stream message");
        if ( read( sockets[0], buf, 1024,0) < 0 )
            perror("reading stream message");
        printf("-->%s\n", buf);
        close(sockets[0]);
    }
}

```

Figure 3 Use of a socketpair

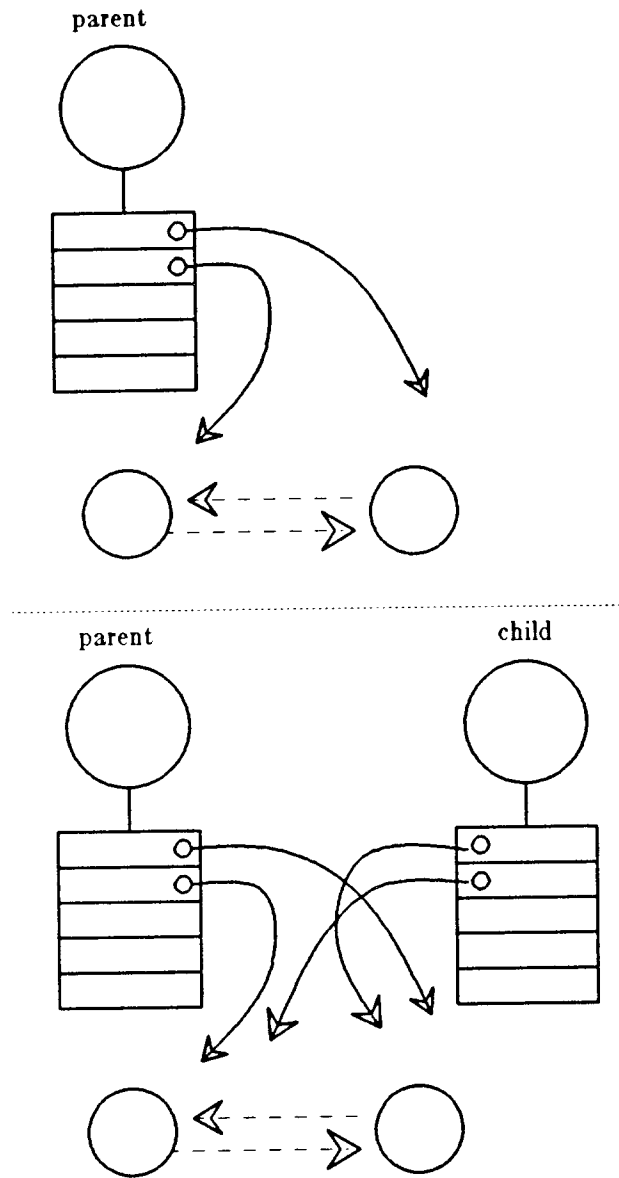


Figure 4 Sharing a socketpair between parent and child

mechanisms require a common ancestor to set up the communication. In other words, we would like to have two processes separately create sockets and then have messages sent between them. This is often the case when providing or using a service in the system. In Berkeley UNIX 4.2BSD one can create individual sockets, give them names and send messages between them.

Sockets created by different programs use names to refer to one another. The space from which a name is drawn is referred to as a *domain*. There are currently two domains for sockets, the UNIX domain (or AF_UNIX, for Address Format UNIX) and the Internet domain (or

AF_INET). In the UNIX domain, a socket is given a path name within the file system name space. An i-node is created for the socket and other processes may then refer to the socket by giving the proper pathname. Names in the internet domain consist of a machine network address and an identifying number, called a port. UNIX domain names, therefore, allow communication between any two processes that work in the same file system. Internet domain names allow communication between machines.

Communication follows some particular "style." Currently, communication is either through a *stream* or by *datagram*. Stream communication implies several things. Communication takes place across a connection between two sockets. The communication is reliable, error-free, and, as in pipes, no message boundaries are kept. Reading from a stream may result in reading the data sent from one or several calls to *write()* or only part of the data from a single call, if there is not enough room for the entire message, or if not all the data from a large message has been transferred. The protocol implementing such a style will retransmit messages received with errors. It will also return error messages if one tries to send a message after the connection has been broken. Datagram communication does not use connections. Each message is addressed individually. If the address is correct, it will generally be received, although this is not guaranteed. Often datagrams are used for requests that require a response from the recipient. If no response arrives in a reasonable amount of time, the request is repeated. The individual datagrams will be kept separate when they are read, that is, message boundaries are preserved.

The difference in performance between the two styles of communication is generally less important than the difference in semantics. The performance gain that one might find in using datagrams must be weighed against the increased complexity of the program, which must now concern itself with lost or out of order messages. If lost messages may simply be ignored, the quantity of traffic may be a consideration; the expense of setting up a connection is best justified by frequent use of the connection. Since the performance of a protocol changes as it is tuned for different situations, it is best to seek the most up-to-date information when making choices for a

program in which performance is crucial.

A protocol is a set of rules and conventions that regulate the transfer of data between participants in the communication. The code that implements a protocol keeps track of the names that are bound to sockets, sets up connections and transfers data between sockets, perhaps sending the data across a network. This code also keeps track of the names that are bound to sockets. It is possible for several protocols, differing only in low level details, to implement the same style of communication within a particular domain. Although it is possible to select which protocol should be used, for most users it is simplest to request the default protocol. This has been done in all the example programs.

One specifies the domain, style and protocol of a socket when it is created. For example, in Figure 5a the call to *socket()* causes the creation of a datagram socket with the default protocol in the UNIX domain.

6. Datagrams in the UNIX Domain

Let us now look at two programs that create sockets separately. The programs in Figures 5a and 5b use datagram communication rather than a stream. The structure used to name UNIX domain sockets is defined in the file `<sys/un.h>`. The definition has also been included in the example for clarity.

Each program creates a socket with a call to *socket()*. These sockets are in the UNIX domain. Once a name has been decided upon it is attached to a socket by the system call *bind()*. The program in Figure 5a uses the name "socket", which it binds to its socket. This name will appear in the working directory of the person running the program. The routines in Figure 5b uses its socket only for sending messages. It does not create a name for the socket because no other process has to refer to it.

Names in the UNIX domain are path names. Like file path names they may be either absolute (e.g. `"/dev/imaginary"`) or relative (e.g. `"socket"`). Because these names are used to

```

#include <sys/types.h>
#include <sys/socket.h>
#include <sys/un.h>

/* In the included file <sys/un.h> a sockaddr_un is defined as follows
 * struct    sockaddr_un {
 *     short  sun_family;
 *     char   sun_path[109];
 * };
 */

#include <stdio.h>

#define NAME "socket"

/* This program creates a UNIX domain datagram socket, binds a name to it,
 * then reads from the socket.
 */

main()
{
    int sock, length;
    struct sockaddr_un name;
    char buf[1024];

    /* Create socket from which to read. */
    sock = socket (AF_UNIX, SOCK_DGRAM, 0);
    if (sock < 0) {
        perror("opening datagram socket");
        exit(-1);
    };

    /* Create name. */
    name.sun_family = AF_UNIX;
    strcpy( name.sun_path, NAME );
    if ( bind( sock, &name, sizeof(struct sockaddr_un) ) ) {
        close(sock);
        perror("binding name to datagram socket");
    }

    printf("socket -->%s0, NAME);
    /* Read from the socket */
    if ( read( sock, buf, 1024, 0 ) < 0 )
        perror("receiving datagram packet");
    printf("-->%s0, buf);
    unlink(NAME);
    close(sock);
}

```

Figure 5a Reading using UNIX domain datagrams

```

#include <sys/types.h>
#include <sys/socket.h>
#include <sys/un.h>
#include <stdio.h>

#define DATA "The sea is calm tonight, the tide is full . . ."

/* Here I send a datagram to a receiver whose name I get from the
 * command line arguments. The form of the command line is
 *      udgramsend pathname
 */

main( argc, argv )
int argc;
char *argv[];
{
    int sock;
    struct sockaddr_un name;

    /* Create socket on which to send. */
    sock = socket (AF_UNIX, SOCK_DGRAM, 0);
    if (sock < 0) {
        perror("opening datagram socket");
        exit(-1);
    };

    /* Construct name of socket to send to. */
    name.sun_family = AF_UNIX;
    strcpy( name.sun_path, argv[1] );
    /* Send message. */
    if (sendto(sock, DATA, sizeof(DATA), 0, &name, sizeof(struct sockaddr_un)) < 0) {
        perror("sending datagram message");
    }
    close(sock);
}

```

Figure 5b Sending a UNIX domain datagrams

allow processes to rendezvous, relative path names can pose difficulties and should be used with care. When a name is bound into the name space, an i-node is allocated. If the i-node is not deallocated, the name will continue to exist even after the bound socket is closed. This can cause subsequent runs of a program to find that a name is unavailable, and directories to fill up with these objects. The names are removed by calling *unlink()*. Names in the UNIX domain are only used for rendezvous. They are not used for message delivery once a connection is established. Therefore, in contrast with the Internet domain, unnamed sockets need not be (and are not) automatically given names when they are connected.

There is no established means of communicating names to interested parties. In the example, the program in Figure 5b gets the name of the socket to which it will send its message through its command line arguments. Once a line of communication has been created one can send the names of additional, perhaps new, sockets over the link. Facilities will have to be built that will make the distribution of names less of a problem than it now is.

7. Datagrams in the Internet Domain

The examples in Figure 6a and 6b are very close to the previous example except that the socket is in the Internet domain. The structure of internet domain names is defined in the file `<netinet/in.h>`. Internet domain names specify a delivery slot, or port, on a particular machine. These names are managed by the system routines that implement a particular protocol. Unlike UNIX domain names, Internet domain names are not entered into the file system and, therefore, they do not have to be unlinked after the named socket has been closed. Routines implementing a particular protocol can run on several different machines. When a message must be sent between machines it is sent to the protocol routine on the destination machine, which interprets the Internet name to decide to which socket the message should be delivered. Several different protocols may be active on the same machine, but, in general, they will not talk to one another. As a result, different protocols may each use the same name, that is a particular `<port, machine address>` pair, to indicate a socket using that protocol. Thus, implicitly, an Internet name is a triple including a protocol as well as the port and machine address.

The protocol for a socket is chosen when the socket is created. The machine address for a socket can be any valid network address of the machine, if it has more than one, or it can be the wildcard value `INADDR_ANY`. The wildcard value is used in the program in Figure 6a. If a machine has several network addresses, it is likely that messages sent to any of the addresses should be deliverable to a socket. This will be the case if the wildcard value has been chosen. Note that even if the wildcard value is chosen, a program sending messages to the named socket

```

#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <stdio.h>

/* In the included file <netinet/in.h> a sockaddr_in is defined as follows:
 * struct sockaddr_in {
 *     short  sin_family;
 *     u_short sin_port;
 *     struct in_addr sin_addr;
 *     char   sin_zero[8];
 * };
 */

/* This program creates a datagram socket, binds a name to it, then reads
 * from the socket. */

main()
{
    int sock, length;
    struct sockaddr_in name;
    char buf[1024];

    /* Create socket from which to read. */
    sock = socket (AF_INET, SOCK_DGRAM, 0);
    if (sock < 0) {
        perror("opening datagram socket");
        exit(-1);
    };

    /* Create name with wildcards. */
    name.sin_family = AF_INET;
    name.sin_addr.s_addr = INADDR_ANY;
    name.sin_port = 0;
    if ( bind (sock, &name, sizeof(name)) ) {
        close(sock);
        perror("binding datagram socket");
    }

    /* Find assigned port value and print it out. */
    length = sizeof(name);
    if ( getsockname (sock, &name, &length) ) {
        close(sock);
        perror("getting socket name");
    }

    printf( "Socket has port #%d0, name.sin_port);
    /* Read from the socket */
    if ( read( sock, buf, 1024, 0) < 0 )
        perror("receiving datagram packet");
    printf("-->%s0, buf);
    close(sock);
}

```

Figure 6a Reading using Internet domain datagrams

```

#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
#include <stdio.h>

#define DATA "The sea is calm tonight, the tide is full . . ."

/* Here I send a datagram to a receiver whose name I get from the
 * command line arguments. The form of the command line is
 *   dgramsend hostname portnumber
 */

main( argc, argv )
int argc;
char *argv[];
{
    int sock;
    struct sockaddr_in name;
    struct hostent * hp, * gethostbyname();

    /* Create socket on which to send. */
    sock = socket (AF_INET, SOCK_DGRAM, 0);
    if (sock < 0) {
        perror("opening datagram socket");
        exit(-1);
    };

    /* Construct name of socket to send to.
     * Gethostbyname() returns a structure including the network
     * address of the specified host. The port number is taken from
     * the command line.
     */
    hp = gethostbyname( argv[1] );
    bcopy ( hp->h_addr, &(name.sin_addr.s_addr), hp->h_length);
    name.sin_family = AF_INET;
    name.sin_port = atoi( argv[2] );
    /* Send message. */
    if (sendto(sock, DATA, sizeof(DATA), 0, &name, sizeof(name)) < 0 ) {
        perror("sending datagram message");
    }
    close(sock);
}

```

Figure 6b Sending an Internet domain datagram

must specify a valid network address. One can be willing to receive from "anywhere," but one cannot send a message "anywhere." The program in Figure 6b is given the destination host name as a command line argument. To determine a network address to which it can send the message, it looks the host address up by the call to *gethostbyname()*. The returned structure includes the

host's network address, which is copied into the structure specifying the destination of the message.

The port number can be thought of as the number of a mailbox, into which the protocol places ones messages. Certain daemons have reserved mailboxes with "well-known" port numbers. These fall in the range from 1 to 1023. Higher numbers are available to users. One does not typically ask for a particular number, since this would lead to collisions. Instead one specifies the port number 0, which causes a free port number to be assigned to the socket. Since Internet names are necessary for message delivery in the Internet domain, names are bound to unnamed sockets during a connect. Note that names are not automatically reported back to the user. After calling *bind()*, asking for port 0, one can call *getsockname()* to discover what port was actually assigned. The routine *getsockname()* will not work for names in the UNIX domain.

8. Connections

To send data between stream sockets (having communication style `SOCK_STREAM`), the sockets must be connected. Figures 7a and 7b show two programs that create such a connection. The program in 7b is relatively simple. To initiate a connection, this program simply calls *connect()*, specifying the name of the socket to which it wishes its socket connected. Provided that the target socket exists and is prepared to handle a connection, the program can begin to send messages. Messages will be delivered in order without message boundaries, as with pipes. The connection is destroyed when either end socket is closed (or soon thereafter). If the process persists in sending messages after the connection is closed, a `SIGPIPE` signal is sent to the process by the operating system. Unless explicit action is taken to handle the signal (see the manual page for *signal* or *sigvec*), the process will terminate and the shell will print the message "broken pipe."

Forming a connection is asymmetrical; one process, like that running the program in Figure 7b, requests a connection with a particular socket, the other process accepts connection requests. Before a connection can be accepted a socket must be created and a name bound to it. This


```

#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
#include <stdio.h>
#define TRUE 1

/* This program creates a socket and then begins an infinite loop.
 * Each time through the loop it accepts a connection and prints out
 * messages from it. When the connection breaks, or a termination message
 * comes through, the program accepts a new connection.
 */

main()
{
    int sock, length;
    struct sockaddr_in server;
    int msgsock;
    char buf[1024];
    int rval;
    int i;

    /* Create socket */
    sock = socket (AF_INET, SOCK_STREAM, 0);
    if (sock < 0) {
        perror("opening stream socket");
        exit(0);
    };

    /* Name socket using wildcards */
    server.sin_family = AF_INET;
    server.sin_addr.s_addr = INADDR_ANY;
    server.sin_port = 0;
    if ( bind (sock, &server, sizeof(server)) ) {
        perror("binding stream socket");
    }

    /* Find out assigned port number and print it out */
    length = sizeof(server);
    if ( getsockname (sock, &server, &length) ) {
        perror("getting socket name");
        exit(0);
    }
    printf("Socket has port %d, ntohs(server.sin_port));

    /* Start accepting connections */
    listen (sock, 5);
    do {
        msgsock = accept(sock, 0, 0);
        do {
            for ( i = 0; i < 1024; i++) buf[i] = ' ';
            if ( (rval = read(msgsock, buf, 1024)) < 0 )
                perror("reading stream message");
            i = 0;

```

```
        if ( rval == 0 ) printf("Ending connection");  
        else printf("--->%s0, buf);  
        } while ( rval != 0 );  
close(msgsock);  
} while ( TRUE );  
}
```

Figure 7a Accepting an Internet domain stream connection

```

#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
#include <stdio.h>

#define M_SOCKET 0x00000001
#define DATA "Half a league, half a league . . ."

/* This program creates a socket and initiates a connection with the
 * socket given in the command line. One message is sent over the
 * connection and then the socket is closed, ending the connection.
 * The form of the command line is
 *   streamwrite hostname portnumber
 */

main( argc, argv )
int argc;
char **argv;
{
    int sock;
    struct sockaddr_in server;
    struct hostent * hp, * gethostbyname();
    char buf[1024];

    /* Create socket */
    sock = socket (AF_INET, SOCK_STREAM, 0);
    if (sock < 0) {
        perror("opening stream socket");
        exit(0);
    };

    /* Connect socket using name specified by command line. */
    server.sin_family = AF_INET;
    hp = gethostbyname(argv[1]);
    bcopy ( hp->h_addr, &(server.sin_addr.s_addr), hp->h_length);
    server.sin_port = htons(atoi(argv[2]));

    if ( connect(sock, &server, sizeof(server)) < 0 ) {
        close(sock);
        perror("connecting stream socket");
        exit(0);
    };

    if ( write(sock, DATA, sizeof(DATA)) < 0)
        perror("writing on stream socket");
    close(sock);
}

```

Figure 7b Initiating an Internet domain stream connection

```

#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netdb.h>
#include <stdio.h>
#define TRUE 1

/* This program uses select() to check for data before reading.
*/

main()
{
    int sock, length;
    struct sockaddr_in server;
    int msgsock;
    char buf[1024];
    int rval;
    int i;
    int ready;
    struct timeval to;

    /* Create socket */
    sock = socket (AF_INET, SOCK_STREAM, 0);
    if (sock < 0) {
        perror("opening stream socket");
        exit(0);
    };

    /* Name socket using wildcards */
    server.sin_family = AF_INET;
    server.sin_addr.s_addr = INADDR_ANY;
    server.sin_port = 0;
    if ( bind (sock, &server, sizeof(server)) ) {
        perror("binding stream socket");
    }

    /* Find out assigned port number and print it out */
    length = sizeof(server);
    if ( getsockname (sock, &server, &length) ) {
        perror("getting socket name");
        exit(0);
    }
    printf("Socket has port %#d0, ntohs(server.sin_port));

    /* Start accepting connections */
    listen (sock, 5);
    do {
        ready = 1 << sock;
        to.tv_sec = 5;
        select(20, &ready, 0, 0, &to);
        if ( ready ) {
            msgsock = accept(sock, 0, 0);
            do {
                for ( i = 0; i < 1024; i++) buf[i] = ' ';
                if ( (rval = read(msgsock, buf, 1024)) < 0 )

```

```
        perror("reading stream message");
        i = 0;
        if ( rval == 0 ) printf("Ending connection");
        else printf("-->%s", buf);
        } while ( rval != 0 );
        close(msgsock);
    }
    else printf("Do something else");
} while ( TRUE );
}
```

Figure 7c Using select() to check for pending connections




Figure 8 Establishing a stream connection

situation is illustrated in the top half of Figure 8. Process 2 has created a socket and bound a name to it. Process 1 has created an unnamed socket. The name bound to process 2's socket is then made known to process 1 and, perhaps to several other potential communicants as well. If there are several possible communicants, this one socket might get several requests for connections. As a result, a new socket is created for each connection. This new socket is the endpoint of communication within this process for this connection. A connection can be destroyed by closing the corresponding socket.

The program in Figure 7a is a rather trivial example of a server. It creates a socket to which it binds a name, which it then advertises. (In this case it prints out the socket number.) The program then calls *listen()* for this socket. Since several clients may attempt to connect more or less simultaneously, a queue of pending connections is maintained in the system address space. *Listen()* prepares the socket to accept connections by initializing the queue. When a connection is requested, it is listed in the queue. If the queue is full, an error status is returned to the requestor. The maximum length of this queue is specified by the second argument of *listen()* (the maximum length may be limited by the system). Once the listen call has been completed, the program enters an infinite loop. On each pass through the loop, a new connection is accepted and removed from the queue, and, hence, a new socket for the connection is created. The bottom half of Figure 8 shows the result of Process 1 connecting with the named socket of Process 2 and Process 2 accepting the connection. After the connection is created, the service, in this case printing out the messages, is performed and the connection socket closed. The *accept()* call will take a pending connection request from the queue, if one is available, or block, waiting for a request. Messages are read from the connection socket. Reads from an active connection will normally block until data is available. The number of bytes read is returned. When a connection is destroyed, the read call returns immediately. The number of bytes returned will be zero.

The program in Figure 7c is a slight variation on the server in Figure 7a. It avoids blocking when there are no pending connection requests by calling *select()* to check for pending requests before calling *accept()*.

The programs in Figures 9a and 9b show a program using stream communication in the UNIX domain. Streams in the UNIX domain can be used for this sort of program in exactly the same way as Internet domain streams, except for the form of the names and the restriction of the connections to a single file system. There are some differences, however, in the functionality of streams in the two domains, notably in the handling of *out-of-band* data (discussed briefly below). These differences are beyond the scope of this paper.

```

#include <sys/types.h>
#include <sys/socket.h>
#include <sys/un.h>
#include <stdio.h>

#define TRUE 1
#define NAME "socket"

/* This program creates a socket in the UNIX domain and binds a name
 * to it. After printing the socket's name it begins a loop.
 * Each time through the loop it accepts a connection and prints out
 * messages from it. When the connection breaks, or a termination message
 * comes through, the program accepts a new connection.
 */

main()
{
    int sock;
    struct sockaddr_un server;
    int msgsock;
    char buf[1024];
    int rval;
    int i;

    /* Create socket */
    sock = socket (AF_UNIX, SOCK_STREAM, 0);
    if (sock < 0) {
        perror("opening stream socket");
        exit(0);
    };

    /* Name socket using wildcards */
    server.sun_family = AF_UNIX;
    strcpy(server.sun_path, NAME);
    if ( bind (sock, &server, sizeof(struct sockaddr_un)) ) {
        perror("binding stream socket");
    }
    printf("Socket has name %s0, server.sun_path);

    /* Start accepting connections */
    listen (sock, 5);
    do {
        msgsock = accept(sock, 0, 0);
        do {
            for ( i = 0; i < 1024; i++) buf[i] = ' ';
            if ( (rval = read(msgsock, buf, 1024)) < 0 )
                perror("reading stream message");
            if ( rval == 0 ) {
                printf("Ending connection0);
            }
            else {
                printf("-->%s0, buf);
            };
        };
    };
}

```



```

        } while ( rval != 0 );
        close(msgsock);
    } while ( TRUE );
    unlink(NAME);
    close(msgsock);
}

```

Figure 9a Accepting a UNIX domain stream connection

```

#include <sys/types.h>
#include <sys/socket.h>
#include <sys/un.h>
#include <stdio.h>

#define DATA "Half a league, half a league . . ."

/* This program connects to the socket named in the command line and
 * sends a one line message to that socket.
 * The form of the command line is
 *   ustreamwrite pathname
 */

main( argc,argv )
int argc;
char *argv[];
{
    int sock;
    struct sockaddr_un server;
    char buf[1024];

    /* Create socket */
    sock = socket (AF_UNIX,SOCK_STREAM,0);
    if (sock<0) {
        perror("opening stream socket");
        exit(0);
    };

    /* Connect socket using name specified by command line. */
    server.sun_family = AF_UNIX;
    strcpy(server.sun_path,argv[1]);

    if ( connect(sock,&server,sizeof(struct sockaddr_un)) < 0 ) {
        close(sock);
        perror("connecting stream socket");
        exit(0);
    };

    if ( write(sock,DATA,sizeof(DATA)) < 0 )
        perror("writing on stream socket");
}

```

Figure 9b Initiating a UNIX domain stream connection

9. Reads, Writes, Recvs, etc.

UNIX 4.2 BSD has several system calls for reading and writing information. The simplest calls are *read()* and *write()*. *Write()* takes as arguments the index of a descriptor, a pointer to a buffer containing the data and the size of the data. The descriptor may indicate either a file or a connected socket. "Connected" can mean either a connected stream socket (as described in Section 8) or a datagram socket for which a *connect()* call has provided a default destination (see the *connect()* manual page). *Read()* also takes a descriptor that indicates either a file or a socket. *Write()* requires a connected socket since no destination is specified in the parameters of the system call. *Read()* can be used for either a connected or an unconnected socket. These calls are, therefore, quite flexible and may be used to write applications that require no assumptions about the source of their input or the destination of their output. There are variations on *read()* and *write()* that allow the source and destination of the input and output to be several separate buffers, while retaining the flexibility to handle both files and sockets. These are *readv()* and *writev()*, for read and write *vector*.

It is sometimes necessary to send high priority data over a connection that may have unread low priority data at the other end. For example, a user interface process may be interpreting commands and sending them on to another process through a stream connection. The user interface may have filled the stream with as yet unprocessed requests when the user types a command to cancel all outstanding requests. Rather than have the high priority data wait to be processed after the low priority data, it is possible to send it as *out-of-band* (OOB) data. The reception of OOB data results in the generation of a SIGURG signal, if this signal has been enabled (see the manual page for *signal* or *sigvec*). See [Leffler 1983] for a more complete description of the OOB mechanism. There are a pair of calls allowing the sending and receiving of OOB information. These are *send()* and *recv()*. These calls require the use of sockets; specifying a file descriptor will result in the return of an error status. These calls also allow *peeking* at data in a stream. That is they allow a process to read data without removing the data

from the stream. One use of this facility is to read ahead in a stream to determine the size of the next item to be read. When not using these options, these calls have the same functions as *read()* and *write()*.

To send datagrams, one must be allowed to specify the destination. The call *sendto()* takes as input a destination address and is therefore used for sending datagrams. The call *recvfrom()* is often used to read datagrams, since this call returns the name of the sender, if it is available. If the identity of the sender does not matter, one may use *read()* or *recv()*.

Finally, there are a pair of calls that allow the sending and receiving of messages from multiple buffers, when the name of the recipient must be specified. These are *sendmsg()* and *recvmsg()*. These calls are actually quite general and have other uses, including, in the UNIX domain, the handing of a socket from one process to another.

The various options for reading and writing are shown in Figure 10, together with their parameters. The parameters for each system call reflect the differences in function of the different calls. In the examples given in this paper, the calls *read()* and *write()* have been used whenever possible.

10. Choices

This paper has presented some of the forms of communication supported by Berkeley UNIX 4.2BSD. These have been presented in an order chosen for ease of presentation. It is useful to review these options emphasizing the factors that make each attractive.

Pipes have the advantage of portability, in that they are supported in all UNIX systems, not just the Berkeley system. They also are relatively simple to use. Socketpairs share this simplicity and have the additional advantage of allowing bidirectional communication. The major shortcoming of these mechanisms is that they require communicating processes to be descendants of a common process. They do not allow intermachine communication.

```

cc= read(descriptor, buf, nbytes)
int cc, descriptor; char *buf; int nbytes;
/* The variable descriptor may be the descriptor of either a file
 * or of a socket.
 */

cc= readv(descriptor, iov, iovcnt)
int cc, descriptor; struct iovec *iov; int iovcnt;
/* An iovec can include several source buffers.
 */

cc = write(descriptor, buf, nbytes)
int cc, descriptor; char *buf; int nbytes;

cc = writev(descriptor, iovec, iovectl)
int cc, descriptor; struct iovec *iovec; int iovectl;

cc= send(socket, msg, len, flags)
int cc, socket; char *msg; int len, flags;
/* The variable socket must be the descriptor of a socket.
 * Flags include OOB and "peeking."
 */

cc= sendto(socket, msg, len, flags, to, tolen)
int cc, socket; char *msg; int len, flags; struct sockaddr *to; int tolen;

cc= sendmsg(socket, msg, flags)
int cc, socket; struct msghdr msg[]; int flags;

cc= recv(socket, buf, len, flags)
int cc, socket; char *buf; int len, flags;

cc= recvfrom(socket, buf, len, flags, from, fromlen)
int cc, socket; char *buf; int len, flags;
struct sockaddr *from; int *fromlen;

cc= recvmsg(socket, msg, flags)
int cc, socket; struct msghdr msg[]; int flags;

```

Figure 10 Varieties of read and write commands

The two naming domains, UNIX and Internet, allow processes with no common ancestor to communicate. Currently, only the Internet domain allows communication between machines. This makes the Internet domain a necessary choice for processes running on separate machines.

The choice between datagrams and stream communication is best made by carefully considering the semantic and performance requirements of the application. Streams can be both advantageous and disadvantageous. One disadvantage is that a process is only allowed a limited

number of open streams, since there are usually only twenty entries available in the open descriptor table. This can cause problems if a single server must talk with a large number of clients. Another is that for delivering a short message the stream setup and teardown time can be unnecessarily large. Weighed against this are the reliability built into the streams. This will often be the deciding factor in favor of streams.

11. What to do Next

Many of the examples presented here can serve as models for multiprocess programs and for programs distributed across several machines. In developing a new multiprocess program, it is often easiest to first write the code to create the processes and communication paths. After this code is debugged, the code specific to the application can be added.

An introduction to the UNIX system and programming using UNIX system calls can be found in [Kernighan and Pike 1984]. Further documentation of the Berkeley UNIX 4.2BSD IPC mechanisms can be found in [Leffler, Fabry & Joy 1983]. More detailed information about particular calls is provided by the UNIX Programmer's Manual [Leffler, Joy & McKusick 1983]. In particular the following manual pages are relevant:

creating and naming sockets: `socket(2)`, `bind(2)`

establishing connections: `listen(2)`, `accept(2)`, `connect(2)`

transferring data: `read(2)`, `write(2)`, `send(2)`, `recv(2)`

addresses: `inet(4F)`

protocols: `tcp(4P)`, `udp(4P)`.

Acknowledgements

I would like to thank Sam Leffler and Mike Karels for their help in understanding the IPC mechanisms and all the people whose comments have helped in writing and improving this report.

References

B.W. Kernighan & R. Pike, 1984,
The UNIX Programming Environment.
Englewood Cliffs, N.J.: Prentice-Hall.

B.W. Kernighan & D.M. Ritchie, 1978,
The C Programming Language,
Englewood Cliffs, N.J.: Prentice-Hall.

S.J. Leffler, R.S. Fabry & W.N. Joy, 1983,
A 4.2BSD Interprocess Communication Primer.
Computer Systems Research Group,
Department of Electrical Engineering and Computer Science,
University of California, Berkeley.

S.J. Leffler, W.N. Joy & M.K. McKusick, 1983,
UNIX Programmer's Manual
Computer Systems Research Group,
Department of Electrical Engineering and Computer Science,
University of California, Berkeley.

Erratum

- In figures 5a, 5b, 9a and 9b, the use of the size `sizeof(struct sockaddr_un)` in some system calls can result in an error. This is due to a known bug in the UNIX kernel that will be fixed in future releases. If this problem occurs in running the example programs on your system, the problem can be fixed by using `sizeof(struct sockaddr_un)-1`.

