

9th Slide Set

Operating Systems

Prof. Dr. Christian Baun

Frankfurt University of Applied Sciences
(1971–2014: Fachhochschule Frankfurt am Main)
Faculty of Computer Science and Engineering
`christianbaun@fb2.fra-uas.de`



Interprocess Communication (IPC)

- Processes do not only carry out read and write operations on data, but also:
 - call each other
 - wait for each other
 - coordinate with each other
 - In short: They must **interact** with each other
- Important questions regarding **interprocess communication (IPC)**:
 - How can a process transmit information to others?
 - How can multiple processes access shared resources?

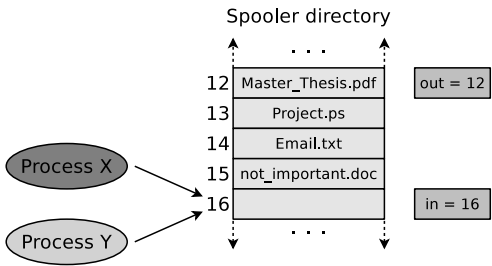

Question: What is the situation here with threads?

- For threads, the same challenges and solutions exist as for interprocess communication with processes
- Only the communication between the threads of a process is no problem because they operate in the same address space

Critical Sections

- If multiple processes run in parallel, the processes consist of...
 - **Uncritical sections:** The processes do not access shared data or carry out only read operations on shared data
 - **Critical sections:** The processes carry out read and write operations on shared data
 - Critical sections must not be processed by multiple processes at the same time
- For processes to be able to access a shared memory (⇒ common data), the operating system must provide **mutual exclusion**


Critical Sections – Example: Print Spooler





- The spooling directory is consistent
 - But the entry of **process Y** was overwritten by **process X** and got lost
- Such a situation is called **race condition**

Signaling

- One way to synchronize processes
- Used to specify an **execution order**
- Example: Section **X** of process P_A must be executed **before** section **Y** of process P_B
 - The `signal` operation signals that process P_A has finished section **X**
 - Perhaps, process P_B must wait for the signal of process P_A



Most Simple Form of Signaling (Busy Waiting)



- The figure shows **busy waiting** at the signal variable `s`
 - The signal variable can be located in a local file, for example
 - Drawback: CPU resources are wasted, because the `wait` operation occupies the processor at regular intervals
- This technique is also called **spinlock** or **polling**

Locking and Unlocking Processes in Linux (2/2)

- Alternative 2: A local file serves as a locking mechanism for mutual exclusion
 - Each process verifies before entering its critical section whether it can open the file exclusively
 - e.g. with the system call `open` or the standard library function `fopen`
 - If this is not the case, it must pause for a certain time (e.g. with the system call `sleep`) and then try again (**busy waiting**).
 - Alternatively, it can pause itself with `sleep` or `pause` and hope that the process that has already opened the file unblocks it with a signal at the end of its critical section (**passive waiting**)

Summary: Difference between Signaling and Blocking

- **Signaling** specifies the execution order
Example: Execute section X of process P_A before section Y of P_B
- **Blocking / Locking** secures critical sections
The execution order of the critical sections of the processes is not specified! It is just ensured that the execution of critical sections does not overlap

Deadlock Detection with Matrices – Example (2/2)

- If process 3 finished execution, it deallocates its resources

Available resource vector = (2 2 2 0)

Request matrix = $\begin{bmatrix} 2 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ - & - & - & - \end{bmatrix}$

- 2 resources of class 1 are available
- 2 resources of class 2 are available
- 2 resources of class 3 are available
- No resources of class 4 are available
- If process 2 finished execution, it deallocates its resources

- Process 1 is blocked, because no free resources of class 4 exist
- **Process 2 is not blocked**

Available resource vector = (4 2 2 1)

Request matrix = $\begin{bmatrix} 2 & 0 & 0 & 1 \\ - & - & - & - \\ - & - & - & - \end{bmatrix}$

- **Process 1 is not blocked** \implies no deadlock in this example

Create a (System V) Shared Memory Segment (in C)

```
1 #include <sys/ipc.h>
2 #include <sys/shm.h>
3 #include <stdio.h>
4 #define MAXMEMSIZE 20
5
6 int main(int argc, char **argv) {
7     int shared_memory_id = 12345;
8     int returncode_shmget;
9
10    // Create shared memory segment or access an existing one
11    // IPC_CREAT = create a shared memory segment, if it does not still exist
12    // 0600 = Access privileges for the new message queue
13    returncode_shmget = shmget(shared_memory_id, MAXMEMSIZE, IPC_CREAT | 0600);
14
15    if (returncode_shmget < 0) {
16        printf("Unable to create the shared memory segment.\n");
17        perror("shmget");
18    } else {
19        printf("The shared memory segment has been created.\n");
20    }
21 }
```

```
$ ipcs -m
----- Shared Memory Segments -----
key      shmid    owner    perms    bytes    nattch   status
0x00003039 56393780 bnc      600      20       0

$ printf "%d\n" 0x00003039          # Convert from hexadecimal to decimal
12345
```


Detach a (System V) Shared Memory Segment (in C)

```

1  #include <sys/types.h>
2  #include <sys/ipc.h>
3  #include <sys/shm.h>
4  #include <stdio.h>
5  #define MAXMEMSIZE 20
6
7  int main(int argc, char **argv) {
8      int shared_memory_id = 12345;
9      int returncode_shmget;
10     int returncode_shmctl;
11     char *sharedmempointer;
12
13     // Create shared memory segment or access an existing one
14     returncode_shmget = shmget(shared_memory_id, MAXMEMSIZE, IPC_CREAT | 0600);
15     ...
16
17     // Attach the shared memory segment
18     sharedmempointer = shmat(returncode_shmget, 0, 0);
19     ...
20
21     // Detach the shared memory segment
22     returncode_shmctl = shmctl(sharedmempointer);
23     if (returncode_shmctl < 0) {
24         printf("Unable to detach the shared memory segment.\n");
25         perror("shmctl");
26     } else {
27         printf("The shared memory segment has been detached.\n");
28     }
29 }
30 }
```

Erase a (System V) Shared Memory Segment (in C)

```
1 #include <sys/types.h>
2 #include <sys/ipc.h>
3 #include <sys/shm.h>
4 #include <stdio.h>
5 #define MAXMEMSIZE 20
6
7 int main(int argc, char **argv) {
8     int shared_memory_id = 12345;
9     int returncode_shmget;
10    int returncode_shmctl;
11    char *sharedmempointer;
12
13    // Create shared memory segment or access an existing one
14    returncode_shmget = shmget(shared_memory_id, MAXMEMSIZE, IPC_CREAT | 0600);
15    ...
16
17    // Erase shared memory segment
18    returncode_shmctl = shmctl(returncode_shmget, IPC_RMID, 0);
19    if (returncode_shmctl == -1) {
20        printf("Unable to erase the shared memory segment.\n");
21        perror("semctl");
22    } else {
23        printf("The shared memory segment has been erased.\n");
24    }
25 }
26 }
```


Create (System V) Message Queues (in C)

```

1 #include <stdlib.h>
2 #include <sys/types.h>
3 #include <sys/ipc.h>
4 #include <stdio.h>
5 #include <sys/msg.h>
6
7 int main(int argc, char **argv) {
8     int returncode_msgget;
9
10     // Create message queue or access an existing one
11     // IPC_CREAT => create a message queue, if it does not still exist
12     // 0600 = Access privileges for the new message queue
13     returncode_msgget = msgget(12345, IPC_CREAT | 0600);
14     if(returncode_msgget < 0) {
15         printf("Unable to create the message queue.\n");
16         exit(1);
17     } else {
18         printf("The message queue 12345 with the ID %i has been created.\n",
19             returncode_msgget);
20     }
21 }

```

```

$ ipcs -q
----- Message Queues -----
key       msqid      owner      perms     used-bytes  messages
0x00003039 98304      bnc        600       0            0

$ printf "%d\n" 0x00003039      # Convert from hexadecimal to decimal
12345

```


Result of writing a Message into a Message Queue

- Before...

```
$ ipcs -q
----- Message Queues -----
key          msqid      owner      perms      used-bytes  messages
0x00003039  98304     bnc        600         0             0
```

- Afterwards...

```
$ ipcs -q
----- Message Queues -----
key          msqid      owner      perms      used-bytes  messages
0x00003039  98304     bnc        600         80            1
```

Pick a Message from a (System V) Message Queue (in C)

```
1 #include <stdlib.h>
2 #include <sys/types.h>
3 #include <sys/ipc.h>
4 #include <stdio.h>
5 #include <sys/msg.h>
6 #include <string.h>           // This header file is required for strcpy()
7 struct msgbuf {              // Template of a buffer for msgsnd and msgrcv
8     long mtype;              // Message type
9     char mtext[80];          // Send buffer
10 } msg;
11
12 int main(int argc, char **argv) {
13     int returncode_msgget, returncode_msgrcv;
14     msg receivebuffer;        // Create a receive buffer
15
16     // Create message queue or access an existing one
17     returncode_msgget = msgget(12345, IPC_CREAT | 0600)
18
19     msg.mtype = 1;            // Pick the first message of type 1
20     // MSG_NOERROR => The message will be truncated when it is too long
21     // IPC_NOWAIT => Do not block the process if no message exists
22     returncode_msgrcv = msgrcv(returncode_msgget, &msg, sizeof(msg.mtext), msg.mtype,
23                               MSG_NOERROR | IPC_NOWAIT);
24     if (returncode_msgrcv < 0) {
25         printf("Unable to pick a message from the message queue.\n");
26         perror("msgrcv");
27     } else {
28         printf("This message was picked from the message queue: %s\n", msg.mtext);
29         printf("The received message is %i characters long.\n", returncode_msgrcv);
30     }
```

Erase a (System V) Message Queue (in C)

```
1 #include <stdlib.h>
2 #include <sys/types.h>
3 #include <sys/ipc.h>
4 #include <stdio.h>
5 #include <sys/msg.h>
6
7 int main(int argc, char **argv) {
8     int returncode_msgget;
9     int returncode_msgctl;
10
11     // Create message queue or access an existing one
12     returncode_msgget = msgget(12345, IPC_CREAT | 0600);
13     ...
14
15     // Erase message queue
16     returncode_msgctl = msgctl(returncode_msgget, IPC_RMID, 0);
17     if (returncode_msgctl < 0) {
18         printf("Unable to erase the message queue with the ID %i.\n", returncode_msgget);
19         perror("msgctl");
20         exit(1);
21     } else {
22         printf("The message queue with the ID %i has been erased.\n", returncode_msgget);
23     }
24     exit(0);
25 }
```

One example of working with System V message queues in Linux can be found on the website of this course

Message Queues in Linux (System V vs. POSIX)

- The functions described so far for working with message queues are part of the **System V** interface
- Some developers prefer the System V API and Others the POSIX API. . . _(`\`)_/_/


C function calls for POSIX message queue specified in the header file `mqqueue.h`

- `mq_open()`: Create a message queue or access an existing one
- `mq_send()`: Write (send) a message into a message queue. Blocking function
- `mq_timedsend()`: Write (send) a message into a message queue. Blocking instruction with a timeout
- `mq_receive()`: Read (receive) a message from a message queue. Blocking instruction
- `mq_timedreceive()`: Read (receive) a message from a message queue. Blocking instruction with a timeout
- `mq_getattr()`: Request the attributes of a message queue. These are: number of messages in the queue, maximum message size, maximum number of messages. . .
- `mq_setattr()`: Modify the attributes of a message queue
- `mq_notify()`: Notify the process as soon as a message is available
- `mq_close()`: Close a message queue
- `mq_unlink()`: Erase a message queue
- POSIX message queues are created In Linux in the folder `/dev/mqueue`

One example of working with POSIX message queues in Linux can be found on the website of this course

Anonymous Pipes (1/2)

- Pipes can be **anonymous pipes** or **named pipes** (see slide 44)
- An **anonymous pipe**...
 - is a buffered unidirectional communication channel between 2 processes
 - If communication in both directions shall be possible at the same time, 2 pipes are necessary – one for each communication direction
 - operates according to the FIFO principle
 - has a limited capacity
 - Pipe = filled \implies the writing process gets blocked
 - Pipe = empty \implies the reading process gets blocked
 - is created with the system call `pipe()`
 - During this process, the kernel of the operating system creates an Inode (\implies slide set 6) and 2 file descriptors (*handles*)
 - Processes access the access identifiers with `read()` and `write()` system calls (or standard library functions) for reading data from or writing data into the pipe



Anonymous Pipes (2/2)

- When child processes are created with `fork()`, the child processes also inherit access to the file descriptors
- **Anonymous pipes** allow process communication only between closely related processes
 - Only processes, which are closely related via `fork()` can communicate with each other via anonymous pipes
 - If the last process, which has access to an anonymous pipe, terminates, the pipe gets erased by the operating system

Overview of the pipes in Linux/UNIX: `lsdf` | `grep pipe`

Anonymous Pipe Example (in C) – Part 1/2

You can monitor the anonymous pipe in Linux/UNIX via `ls -l | grep -n -P | grep <PID>` and inside the directory `/proc/<PID>/fd`

```
1 #include <stdio.h>
2 #include <unistd.h>
3 #include <stdlib.h>
4
5 void main() {
6     int pid_of_child;
7     // Create handles for the pipe to read (testpipe[0]) and write (testpipe[1])
8     int testpipe[2];
9
10    // Create anonymous pipe testpipe
11    if (pipe(testpipe) < 0) {
12        printf("Unable to create the anonymous pipe.\n");
13        // Terminate process
14        exit(1);
15    } else {
16        printf("Created the anonymous pipe testpipe.\n");
17    }
18
19    // Create a child process
20    pid_of_child = fork();
21
22    if (pid_of_child < 0) {
23        perror("Unable to create the child process!\n");
24        // Terminate process
25        exit(1);
26    }
```

Anonymous Pipe Example (in C) – Part 2/2

```

27 // Parent process
28 if (pid_of_child > 0) {
29     printf("Parent process: PID: %i\n", getpid());
30     // Block the read channel of the anonymous pipe testpipe
31     close(testpipe[0]);
32     char message[] = "Testnachricht";
33     // Write the message into the write channel of the anonymous pipe
34     write(testpipe[1], &message, sizeof(message));
35 }
36
37 // Child process
38 if (pid_of_child == 0) {
39     printf("Child process: PID: %i\n", getpid());
40     // Block the write channel of the anonymous pipe testpipe
41     close(testpipe[1]);
42     // Create a receive buffer (80 bytes capacity)
43     char puffer[80];
44     // Read the message from the read channel of the anonymous pipe
45     read(testpipe[0], puffer, sizeof(puffer));
46     printf("Received: %s\n", puffer);
47 }
48 }

```

```

$ gcc anonymous_pipe_example.c -o anonymous_pipe_example
$ ./anonymous_pipe_example
Created the anonymous pipe testpipe.
Parent process: PID: 394769
Child process: PID: 394770
Received: Testnachricht

```

Named Pipes

- Processes, which are not closely related with each other, can communicate via **named pipes**
 - These pipes can be accessed by using their names
 - They are created in C by: `mkfifo("<pathname>", <permissions>)`
 - Any process, which knows the name of a pipe, can use the name to access the pipe and communicate with other processes
- The operating system ensures **mutual exclusion**
 - At any time, only a single process can access a pipe
- Named pipes are not erased automatically by the operating system (unlike anonymous pipes)

Named Pipe Example (in C) – Part 1/4

```

1 #include <stdio.h>
2 #include <unistd.h>
3 #include <stdlib.h>
4 #include <fcntl.h>
5 #include <sys/stat.h>
6
7 void main() {
8     int pid_of_child;
9
10    // Create named pipe
11    if (mkfifo("testfifo",0666) < 0) {
12        printf("Unable to create the named pipe.\n");
13        exit(1);
14    } else {
15        printf("Created the named pipe testfifo.\n");
16    }
17
18    // Create a child process
19    pid_of_child = fork();
20
21    if (pid_of_child < 0) {
22        perror("Unable to create the child process!\n");
23        exit(1);
24    }

```

The function call creates a file system entry named `testfifo` in the current directory. The first letter in the output of the `ls` command shows that `testfifo` is a named pipe.

```
$ ls -la testfifo
```

```
prw-r--r-- 1 bnc bnc 0 1. Feb 10:15 testfifo
```

Named Pipe Example (in C) – Part 2/4

```
25 // Parent process
26 if (pid_of_child > 0) {
27     printf("Parent process: PID: %i\n", getpid());
28
29     // Create the file descriptor (handle) for the pipe
30     int fd;
31
32     // Specify the message to be transferred
33     char message[] = "Testnachricht";
34
35     // Open the named pipe for writing
36     fd = open("testfifo", O_WRONLY);
37
38     // Write the message into the pipe
39     write(fd, &message, sizeof(message));
40
41     // Close the named pipe
42     close(fd);
43 }
```

Named Pipe Example (in C) – Part 3/4

```

44 // Child process
45 if (pid_of_child == 0) {
46     printf("Child process: PID: %i\n", getpid());
47
48     // Create the file descriptor (handle) for the pipe
49     int fd;
50     // Create a receive buffer
51     char puffer[80];
52
53     // Open the named pipe for reading
54     fd = open("testfifo", O_RDONLY);
55
56     // Read the message from the pipe
57     read(fd, puffer, sizeof(puffer));
58     printf("Received: %s\n", puffer);
59
60     // Close the named pipe
61     close(fd);
62
63     // Erase the named pipe
64     if (unlink("testfifo") < 0) {
65         printf("Unable to erase the named pipe.\n");
66         exit(1);
67     } else {
68         printf("The named pipe has been erased.\n");
69     }
70 }
71 }

```


Named Pipe Example (in C) – Part 4/4

```
$ gcc named_pipe_example.c -o named_pipe_example
$ ./named_pipe_example
Created the named pipe testfifo.
Parent process: PID: 395415
Child process: PID: 395416
Received: Testnachricht
The named pipe has been erased.
```

You can monitor the named pipe in Linux/UNIX via `ls -l | grep <PID>` and inside the directory `/proc/<PID>/fd`

Sockets

- Full duplex-ready alternative to pipes and shared memory
 - Allow interprocess communication in distributed systems
- An user process can request a socket from the operating system and afterwards send and receive data via the socket
 - The operating system maintains all used sockets and the related connection information



- Ports are used for the communication via sockets
 - Port numbers are randomly assigned during connection establishment
 - Port numbers are assigned randomly by the operating system
 - Exceptions are port numbers of well-known applications, such as HTTP (80) SMTP (25), Telnet (23), SSH (22), FTP (21),...
- Sockets can be used in a blocking (synchronous) and non-blocking (asynchronous) way

Different Types of Sockets

- **Connectionless sockets (= datagram sockets)**
 - Use the Transport Layer protocol UDP
 - Advantage: Better data rate as with TCP
 - Reason: Lesser overhead for the protocol
 - Drawback: Segments may arrive in wrong sequence or may get lost
- **Connection-oriented sockets (= stream sockets)**
 - Use the Transport Layer protocol TCP
 - Advantage: Better reliability
 - Segments cannot get lost
 - Segments always arrive in the correct sequence
 - Drawback: Lower data rate as with UDP
 - Reason: More overhead for the protocol


Using Sockets

- Almost all major operating systems support sockets
 - Advantage: Better portability of applications
- Functions for communication via sockets:
 - Creating a Socket:
`socket()`
 - Binding a socket to a port number and making it ready to receive data:
`bind()`, `listen()`, `accept()` and `connect()`
 - Sending/receiving messages via the socket:
`send()`, `sendto()`, `recv()` and `recvfrom()`
 - Closing eines Socket:
`shutdown()` or `close()`

Overview of the sockets in Linux/UNIX: `netstat -n` or `lsof | grep socket`

Examples of Interprocess communication via sockets (TCP and UDP) in Linux can be found on the website of this course

Connection-less Communication via Sockets – UDP



• Client

- Create socket (`socket`)
- Send (`sendto`) and receive data (`recvfrom`)
- Close socket (`close`)

• Server

- Create socket (`socket`)
- Bind socket to a port (`bind`)
- Send (`sendto`) and receive data (`recvfrom`)
- Close socket (`close`)


Connection-oriented Communication via Sockets – TCP

• Client

- Create socket (`socket`)
- Connect client with server socket (`connect`)
- Send (`send`) and receive data (`recv`)
- Close socket (`close`)

• Server

- Create socket (`socket`)
- Bind socket to a port (`bind`)
- Make socket ready to receive (`listen`)
 - Set up a queue for connection requests. Specifies the number of connection requests, which can be stored in the queue
- Server accepts connections (`accept`)
 - Fetch the first connection request from the queue
- Send (`send`) and receive data (`recv`)
- Close socket (`close`)




Sockets via UDP – Example (Server)

```

1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <string.h>
4 #include <sys/socket.h>
5 #include <netinet/in.h>
6 #include <unistd.h>
7 #include <arpa/inet.h>
8
9 int main(int argc, char *argv[]) {
10     int sd, adresse_laenge;
11     char puffer[1024] = { 0 };
12     struct sockaddr_in adresse, client_adresse;
13     memset(&adresse, 0, sizeof(adresse));
14     memset(&client_adresse, 0, sizeof(client_adresse));
15     adresse.sin_family = AF_INET;
16     adresse.sin_addr.s_addr = INADDR_ANY;
17     adresse.sin_port = htons(atoi(argv[1]));
18
19     sd = socket(AF_INET, SOCK_DGRAM, 0);
20     bind(sd, (struct sockaddr *) &adresse, sizeof(adresse));
21     adresse_laenge = sizeof(client_adresse);
22     rcvfrom(sd, (char *)puffer, sizeof(puffer), 0,
23             (struct sockaddr *) &client_adresse, &adresse_laenge);
24     printf("Empfangene Nachricht: %s\n", puffer);
25     char antwort[] = "Server: Nachricht empfangen.\n";
26     sendto(sd, (const char *)antwort, sizeof(antwort), 0,
27            (struct sockaddr *) &client_adresse, adresse_laenge);
28     close(sd);
29     exit(0);
30 }

```



```

$ gcc udp_server.c -o udp_server
$ ./udp_server 50002


```

Sockets via UDP – Example (Client)

```

1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <string.h>
4 #include <sys/socket.h>
5 #include <netinet/in.h>
6 #include <unistd.h>
7 #include <arpa/inet.h>
8
9 int main(int argc, char *argv[]) {
10     int sd, adresse_laenge;
11     char puffer[1024] = { 0 };
12     struct sockaddr_in adresse;
13     memset(&adresse, 0, sizeof(adresse));
14     adresse.sin_family = AF_INET;
15     adresse.sin_port = htons(atoi(argv[2]));
16     adresse.sin_addr.s_addr = inet_addr(argv[1]);
17
18     sd = socket(AF_INET, SOCK_DGRAM, 0);
19     printf("Bitte Nachricht eingeben: ");
20     fgets(puffer, sizeof(puffer), stdin);
21     adresse_laenge = sizeof(adresse);
22     sendto(sd, (const char *)puffer, strlen(puffer), 0,
23           (struct sockaddr *) &adresse, adresse_laenge);
24     memset(puffer, 0, sizeof(puffer));
25     rcvfrom(sd, (char *)puffer, sizeof(puffer), 0,
26            (struct sockaddr *) &adresse, &adresse_laenge);
27     printf("%s\n", puffer);
28     close(sd);
29     exit(0);
30 }

```



```


$ gcc udp_client.c -o udp_client
$ ./udp_client 127.0.0.1 50002
Bitte Nachricht eingeben: Test
Server: Nachricht empfangen.

```

```

$ ./udp_server 50002
Empfangene Nachricht: Test

```




Sockets via TCP – Example (Server)

```

1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <string.h>
4  #include <sys/socket.h>
5  #include <netinet/in.h>
6  #include <unistd.h>
7  #include <arpa/inet.h>
8
9  int main(int argc, char *argv[]) {
10     int sd, fd, adresse_laenge;
11     char puffer[1024] = { 0 };
12     struct sockaddr_in adresse;
13     memset(&adresse, 0, sizeof(adresse));
14     adresse.sin_family = AF_INET;
15     adresse.sin_addr.s_addr = INADDR_ANY;
16     adresse.sin_port = htons(atoi(argv[1]));
17
18     sd = socket(AF_INET, SOCK_STREAM, 0);
19     bind(sd, (struct sockaddr *) &adresse, sizeof(adresse));
20     listen(sd, 5);
21     adresse_laenge = sizeof(adresse);
22     fd = accept(sd, (struct sockaddr *) &adresse, &adresse_laenge);
23     read(fd, puffer, sizeof(puffer));
24     printf("Empfangene Nachricht: %s\n", puffer);
25     char antwort[]="Server: Nachricht empfangen.\n";
26     write(fd, antwort, sizeof(antwort));
27     close(fd);
28     close(sd);
29     exit(0);
30 }

```



```

$ gcc tcp_server.c -o tcp_server
$ ./tcp_server 50003

```





Sockets via TCP – Example (Client)

```

1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <string.h>
4 #include <sys/socket.h>
5 #include <netinet/in.h>
6 #include <unistd.h>
7 #include <arpa/inet.h>
8
9 int main(int argc, char *argv[]) {
10     int sd;
11     char puffer[1024] = { 0 };
12     struct sockaddr_in adresse;
13     memset(&adresse, 0, sizeof(adresse));
14     adresse.sin_family = AF_INET;
15     adresse.sin_port = htons(atoi(argv[2]));
16     adresse.sin_addr.s_addr = inet_addr(argv[1]);
17
18     sd = socket(AF_INET, SOCK_STREAM, 0);
19     connect(sd, (struct sockaddr *) &adresse, sizeof(adresse));
20
21     printf("Bitte Nachricht eingeben: ");
22     fgets(puffer, sizeof(puffer), stdin);
23     write(sd, puffer, strlen(puffer));
24     memset(puffer, 0, sizeof(puffer));
25     read(sd, puffer, sizeof(puffer));
26     printf("%s\n", puffer);
27
28     close(sd);
29     exit(0);
30 }

```



```

$ gcc tcp_client.c -o tcp_client
$ ./tcp_client 127.0.0.1 50003
Bitte Nachricht eingeben: Test
Server: Nachricht empfangen.

```

```



$ ./tcp_server 50003
Empfangene Nachricht: Test

```






A GRAPHIC EXAMPLE OF THE PRODUCER/CONSUMER PROBLEM


Michael Vigneau



The consumer must wait for producer to produce before it can consume...



If the consumer is busy (can't consume), the producer must wait, if the buffer is full, for the consumer to start consuming again. The processes are now _____ (fill in)



Producer/Consumer Example (2/3)

- 3 semaphores are used to synchronize access to the buffer
 - empty
 - filled
 - mutex
- The semaphores filled and empty are used in opposite to each other
 - empty counts the number of empty locations in the buffer and its value is reduced by the producer (P operation) and raised by the consumer (V operation)
 - $\text{empty} = 0 \implies$ puffer is completely filled \implies producer is blocked
 - filled counts the number of data packets (occupied locations) in the buffer and its value is raised by the producer (V operation) and reduced by the consumer (P operation)
 - $\text{filled} = 0 \implies$ puffer is empty \implies consumer is blocked
- The semaphore mutex is used to ensure for the mutual exclusion

Binary Semaphores

- **Binary semaphores** are initialized with value 1 and ensure that 2 or more processes cannot simultaneously enter their critical sections
- Example: The semaphore mutex from the producer/consumer example

Simple Semaphore Example (in C) – Part 4/5

```

79 // Warten auf die Beendigung des Kindprozesses
80 wait(NULL);
81
82 printf("\n");
83
84 // Semaphorgruppe 12345 entfernen
85 returncode_semctl = semctl(returncode_semget1, 0, IPC_RMID, 0);
86 if (returncode_semctl < 0) {
87     printf("Die Semaphorgruppe %i konnte nicht entfernt werden.\n", returncode_semget1);
88     exit(1);
89 } else {
90     printf("Die Semaphorgruppe mit ID %i und Key %i wurde entfernt.\n", returncode_semget1, sem_key1);
91 }
92
93 // Semaphorgruppe 54321 entfernen
94 returncode_semctl = semctl(returncode_semget2, 0, IPC_RMID, 0);
95 if (returncode_semctl < 0) {
96     printf("Die Semaphorgruppe %i konnte nicht entfernt werden.\n", returncode_semget2);
97     exit(1);
98 } else {
99     printf("Die Semaphorgruppe mit ID %i und Key %i wurde entfernt.\n", returncode_semget2, sem_key2);
100 }
101
102 exit(0);
103 }

```

One example of working with semaphores in Linux can be found on the website of this course

Monitor and erase IPC Objects

- Information about existing (**System V**) shared memory segments, (**System V**) message queues and (**System V**) semaphores provides the command `ipcs`
- The easiest way to erase such shared memory segments, message queues and semaphores from the command line is the command `ipcrm`

```
ipcrm [-m shmid] [-q msgqid] [-s semid]  
      [-M shmkey] [-Q msgkey] [-S semkey]
```

- **POSIX** memory segments and **POSIX** semaphores can be inspected and manually erased in the directory `/dev/shm`
- **POSIX** message queues can be inspected and manually erased in the directory `/dev/mqueue`