

جامعة الأنبار

كلية علوم الحاسوب وتكنولوجيا
المعلومات

قسم أنظمة شبكات الحاسوب

المرحلة الرابعة

Operating System

التدريسي: أ.م.د. عمر منذر حسين

1) Process Creation

- During the course of execution, a process may create several new processes.
- The creating process is called a **parent process**, and the new processes are called the **children** of that process.
- Each of these new processes may in turn create other processes, forming a **tree of processes**.
- Most operating systems (including UNIX, Linux, and Windows) identify processes according to a unique **process identifier (or pid)**, which is typically an integer number.
- In general, when a process creates a child process, that child process will need certain resources to accomplish its task.
- A child process may be able to obtain its resources from:
 - the operating system, or
 - it may be constrained to a subset of the resources of the parent process. The parent may have to partition its resources among its children, or it may be able to share some resources (such as memory or files) among several of its children.
- Restricting a child process to a subset of the parent's resources prevents any process from overloading the system by creating too many child processes.
- When a process creates a new process, two possibilities for execution exist:
 1. The parent continues to execute concurrently with its children.
 2. The parent waits until some or all of its children have terminated.

2) Process Termination

- A process terminates when it finishes executing its final statement and asks the operating system to delete it by using the `exit()` system call.
- At that point, the process may return a status value (typically an integer) to its parent process (via the `wait()` system call).
- All the resources of the process are deallocated by the operating system.
- Termination can occur in other circumstances as well.

- A process can cause the termination of another process via an appropriate system call (for example, `TerminateProcess()` in Windows).
- Usually, such a system call can be invoked only by the parent of the process that is to be terminated.
- A parent may terminate the execution of one of its children for a variety of reasons, such as these:
 1. The child has exceeded its usage of some of the resources that it has been allocated.
 2. The task assigned to the child is no longer required.
 3. The parent is exiting, and the operating system does not allow a child to continue if its parent terminates.
- **Some systems do not allow a child to exist if its parent has terminated. In such systems, if a process terminates (either normally or abnormally), then all its children must also be terminated.**
- This phenomenon, referred to as **cascading termination**, is normally initiated by the operating system.

Interprocess Communication

- Processes executing concurrently in the operating system may be either **independent processes or cooperating processes**.
- A process is *independent* if it cannot affect or be affected by the other processes executing in the system.
- Any process that does not share data with any other process is independent.
- A process is *cooperating* if it can affect or be affected by the other processes executing in the system.
- Clearly, any process that shares data with other processes is a cooperating process.
- There are several reasons for providing an environment that allows process cooperation:
 - **Information sharing • Computation speedup • Modularity • Convenience.**
- Cooperating processes require an **interprocess communication (IPC) mechanism** that will allow them to exchange data and information.
- There are two fundamental models of inter-process communication:
 - **shared memory** and **message passing**.
- In the **shared-memory model**, a region of memory that is shared by cooperating processes is established.

- Processes can then exchange information by reading and writing data to the shared region.
- In the message-passing model, communication takes place by means of messages exchanged between the cooperating processes.
- The two communications models are contrasted in Figure 3.12.

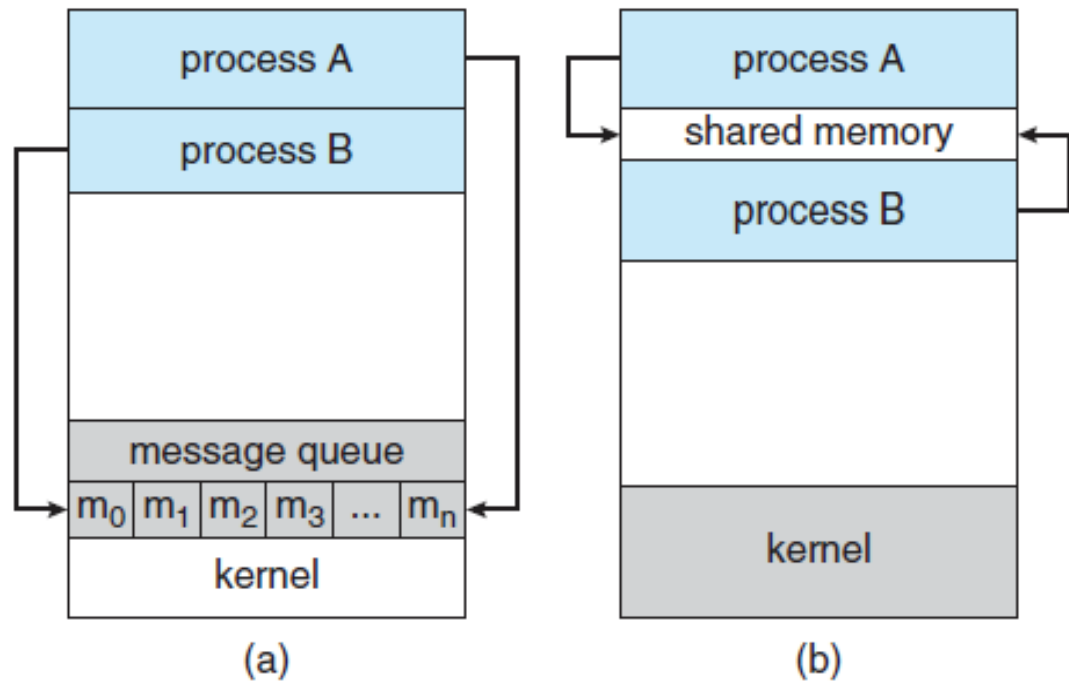
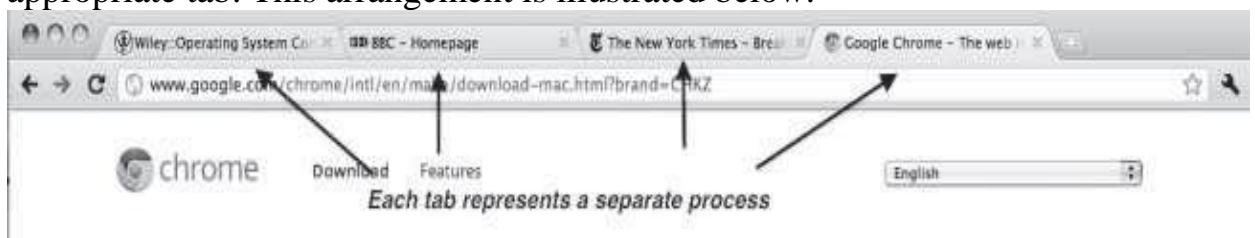


Figure 3.12 Communications models. (a) Message passing. (b) Shared memory.

- Shared memory can be faster than message passing, since message-passing systems are typically implemented using system calls and thus require the more time-consuming task of kernel intervention.
- Most contemporary web browsers provide tabbed browsing, which allows a single instance of a web browser application to open several websites at the same time, with each site in a separate tab.
- To switch between the different sites, a user need only click on the appropriate tab. This arrangement is illustrated below.



- A problem with this approach is that if a web application in any tab crashes, the entire
- Google's Chrome web browser was designed to address this issue by using a multiprocess architecture. Chrome identifies three different types of processes: **browser**, **renderers**, and **plug-ins**.
- **The browser process** is responsible for managing the user interface as well as disk and network I/O.
- A new browser process is created when Chrome is started. Only one browser process is created.
- **Renderer processes** contain logic for rendering web pages.
- Thus, they contain the logic for handling HTML, Javascript, images, and so forth.
- A new renderer process is created for each website opened in a new tab, and so **several renderer processes may be active at the same time**.
- process—including all other tabs displaying additional websites—crashes as well.
- **A plug-in process** is created for each type of plug-in (such as Flash or QuickTime) in use.
- Plug-in processes contain the code for the plug-in as well as additional code that enables the plug-in to communicate with associated renderer processes and the browser process.
- The advantage of the multiprocess approach is that: **websites run in isolation from one another**.
- If one website crashes, only its renderer process is affected; all other processes remain unharmed.
- Furthermore, **renderer processes run in a sandbox**, which means that access to disk and network I/O is restricted, minimizing the effects of any security exploits.

Shared-Memory Systems

- The communicated processes can exchange information by reading and writing data in the shared areas.
- The form of the data and the location are determined by these processes and are not under the operating system's control.

- Recall that, normally, the operating system tries to prevent one process from accessing another process's memory.
- Shared memory requires that two or more processes agree to remove this restriction.
- The processes are also responsible for ensuring that they are not writing to the same location simultaneously.
- To illustrate the concept of cooperating processes, let's consider the producer-consumer problem, which is a common paradigm for cooperating processes.
- A **producer** process produces information that is consumed by a **consumer** process.
- For example, a compiler may produce assembly code that is consumed by an assembler. The assembler, in turn, may produce object modules that are consumed by the loader.

Message-Passing Systems

- Message passing provides a mechanism to allow processes to communicate and to synchronize their actions without sharing the same address space.
- It is particularly useful in a distributed environment, where the communicating processes may reside on different computers connected by a network.
- For example, an Internet chat program could be designed so that chat participants communicate with one another by exchanging messages.
- A message-passing facility provides at least two operations:

send(message) receive(message)

- If processes *P* and *Q* want to communicate, they must send messages to and receive messages from each other.
- A **communication link** must exist between them.
- This link can be implemented in a variety of ways.
- Here are several methods for **logically implementing** a link and the send()/receive() operations:

- **Direct or indirect communication**
- **Synchronous or asynchronous communication**
- **Automatic or explicit buffering**
- We look at issues related to each of these features next.

1- Naming

- Processes that want to communicate must have a way to refer to each other.
- They can use either direct or indirect communication.
- Under **direct communication**, each process that wants to communicate must explicitly name the recipient or sender of the communication.
- In this scheme, the `send()` and `receive()` primitives are defined as:
 - `send(P, message)`—Send a message to process P.
 - `receive(Q, message)`—Receive a message from process Q.
- A communication link in this scheme has the following properties:
 - A link is established automatically between every pair of processes that want to communicate. The processes need to know only each other's identity to communicate.
 - A link is associated with exactly two processes.
 - Between each pair of processes, there exists exactly one link.
- This scheme exhibits ***symmetry in addressing***; that is, both the sender process and the receiver process must name the other to communicate.
- A variant of this scheme employs ***asymmetry in addressing***. Here, only the sender names the recipient; the recipient is not required to name the sender.
 - In this scheme, the `send()` and `receive()` primitives are defined as follows:
 - **`send(P, message)`— Send a message to process P.**

- **receive(id, message)**— **Receive a message from any process. The variable id is set to the name of the process with which communication has taken place.**

2- Synchronization

- Communication between processes takes place through calls to send() and receive() primitives.
- There are different **design options for implementing each primitive.**
- Message passing may be either **blocking** or **nonblocking**— also known as **synchronous** and **asynchronous**.
 - **Blocking send.** The sending process is blocked until the message is received by the receiving process or by the mailbox.
 - **Nonblocking send.** The sending process sends the message and resumes operation.
 - **Blocking receive.** The receiver blocks until a message is available.
 - **Nonblocking receive.** The receiver retrieves either a valid message or a null.

3- Buffering

- Whether communication is direct or indirect, messages exchanged by communicating processes reside in a temporary queue.
- Basically, such queues can be implemented in three ways:
- **Zero capacity.** The queue has a maximum length of zero; thus, the link cannot have any messages waiting in it. In this case, the sender must block until the recipient receives the message.
- **Bounded capacity.** The queue has finite length **n**; thus, at most **n** messages can reside in it.
- If the queue is not full when a new message is sent, the message is placed in the queue, and the sender can continue execution without waiting.

- The link's capacity is finite, however. If the link is full, the sender must block until space is available in the queue.
- **Unbounded capacity.** The queue's length is potentially infinite; thus, any number of messages can wait in it. The sender never blocks.
- The zero-capacity case is sometimes referred to as a message system with no buffering. The other cases are referred to as systems with automatic buffering.
- An Example: Windows
- Windows provides support for multiple operating environments, or **subsystems**.
- Application programs communicate with these subsystems via a message-passing mechanism.
- Thus, application programs can be considered clients of a subsystem server.
- The message-passing facility in Windows is called the **advanced local procedure call (ALPC)** facility.
- It is used for communication between two processes on the same machine.
- Windows uses a **port object** to establish and maintain a connection between two processes.
- Windows uses two types of ports: **connection ports** and **communication ports**.
- Server processes publish connection-port objects that are visible to all processes.
- When a client wants services from a subsystem, it opens a handle to the server's connection-port object and sends a connection request to that port.
- The server then creates a channel and returns a handle to the client.
- The channel consists of a pair of private communication ports: one for client—server messages, the other for server—client messages.
- Additionally, communication channels support a callback mechanism that allows the client and server to accept requests when they would normally be expecting a reply.

- When an ALPC channel is created, one of three message-passing techniques is chosen:
- **1.** For small messages (up to 256 bytes), the port's message queue is used as intermediate storage, and the messages are copied from one process to the other.
- **2.** Larger messages must be passed through a section object, which is a region of shared memory associated with the channel.
- **3.** When the amount of data is too large to fit into a section object, an API is available that allows server processes to read and write directly into the address space of a client.
- The structure of advanced local procedure calls in Windows is shown in Figure 3.19.

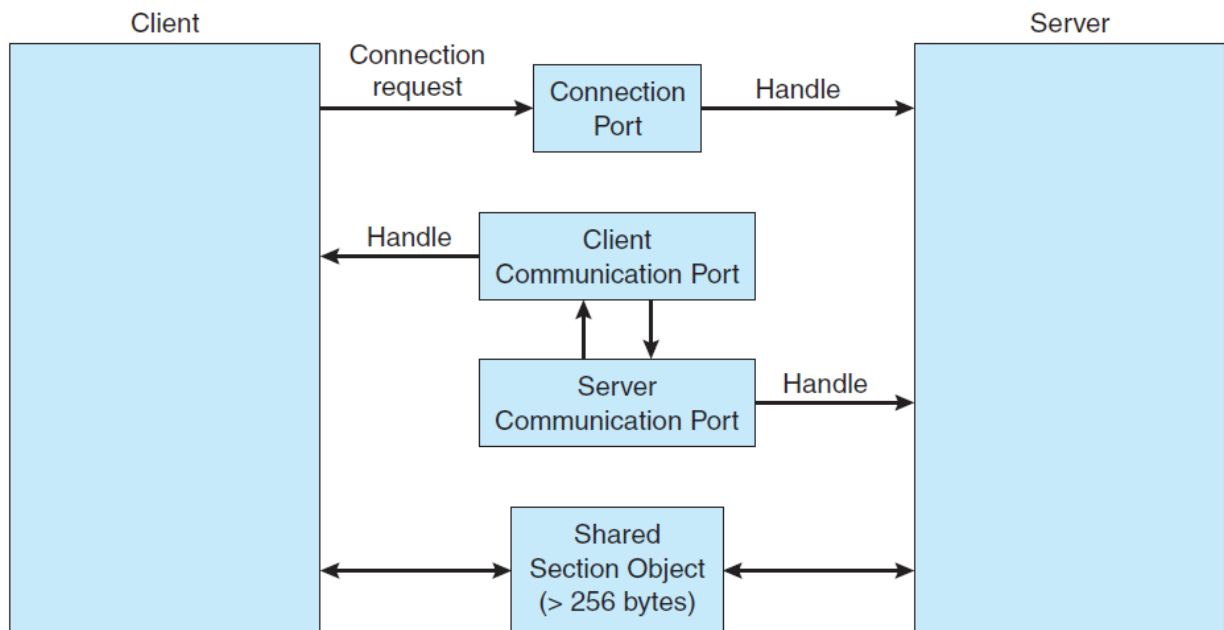


Figure 3.19 Advanced local procedure calls in Windows.