

# *Concurrency and Parallelism*

## *Master 1 International*

---



**Andrea G. B. Tettamanzi**

Université de Nice Sophia Antipolis

Département Informatique

[andrea.tettamanzi@unice.fr](mailto:andrea.tettamanzi@unice.fr)

# *Web Page*

<http://www.i3s.unice.fr/~tettaman/Classes/ConcPar/>

## *Lecture 1*

# **Processes and Threads**

# *Introduction to Threads*

## Basic idea

We build virtual processors in software, on top of physical processors:


- Processor: Provides a set of instructions along with the capability of automatically executing a series of those instructions.
- Thread: A minimal software processor in whose context a series of instructions can be executed. Saving a thread context implies stopping the current execution and saving all the data needed to continue the execution at a later stage.
- Process: A software processor in whose context one or more threads may be executed. Executing a thread, means executing a series of instructions in the context of that thread.

# Context Switching

## Contexts

- Processor context: The minimal collection of values stored in the registers of a processor used for the execution of a series of instructions (e.g., stack pointer, addressing registers, program counter).
- Thread context: The minimal collection of values stored in registers and memory, used for the execution of a series of instructions (i.e., processor context, state).
- Process context: The minimal collection of values stored in registers and memory, used for the execution of a thread (i.e., thread context, but now also at least MMU register values).

# Thread Usage in Nondistributed Systems



## Context switching as the result of IPC

# *Context Switching : Observations*

- Threads share the same address space. Thread context switching can be done entirely independent of the operating system.
- Process switching is generally more expensive as it involves getting the OS in the loop, i.e., trapping to the kernel.
- Creating and destroying threads is much cheaper than doing so for processes.

# *Threads and Operating Systems*


- Main issue : Should an OS kernel provide threads, or should they be implemented as user-level packages?
- User-space solution
  - All operations can be completely handled within a single process  $\Rightarrow$  implementations can be extremely efficient.
  - All services provided by the kernel are done on behalf of the process in which a thread resides  $\Rightarrow$  if the kernel decides to block a thread, the entire process will be blocked.
  - Threads are used when there are lots of external events: threads block on a per-event basis  $\Rightarrow$  if the kernel can't distinguish threads, how can it support signaling events to them?



# *Threads and Operating Systems*

- Kernel solution : The whole idea is to have the kernel contain the implementation of a thread package. This means that all operations return as system calls
  - Operations that block a thread are no longer a problem: the kernel schedules another available thread within the same process.
  - Handling external events is simple: the kernel (which catches all events) schedules the thread associated with the event.
  - The big problem is the loss of efficiency due to the fact that each thread operation requires a trap to the kernel.
- Conclusion: Try to mix user-level and kernel-level threads into a single concept.

# Solaris Threads



Combining kernel-level lightweight processes and user-level threads.

# *Solaris Thread Operation*

- User-level thread does system call  $\Rightarrow$  the LWP that is executing that thread, blocks. The thread remains bound to the LWP.
- The kernel can schedule another LWP having a runnable thread bound to it. Note: this thread can switch to any other runnable thread currently in user space.
- A thread calls a blocking user-level operation  $\Rightarrow$  do context switch to a runnable thread, (then bound to the same LWP).
- When there are no threads to schedule, an LWP may remain idle, and may even be removed (destroyed) by the kernel.
- **Note:** This concept has been virtually abandoned – it's just either user-level or kernel-level threads.

# *Threads in Distributed Systems*

## **Multithreaded Web client**

Hiding network latencies:

- Web browser scans an incoming HTML page, and finds that more files need to be fetched.
- Each file is fetched by a separate thread, each doing a (blocking) HTTP request.
- As files come in, the browser displays them.

## **Multiple request-response calls to other machines (RPC)**

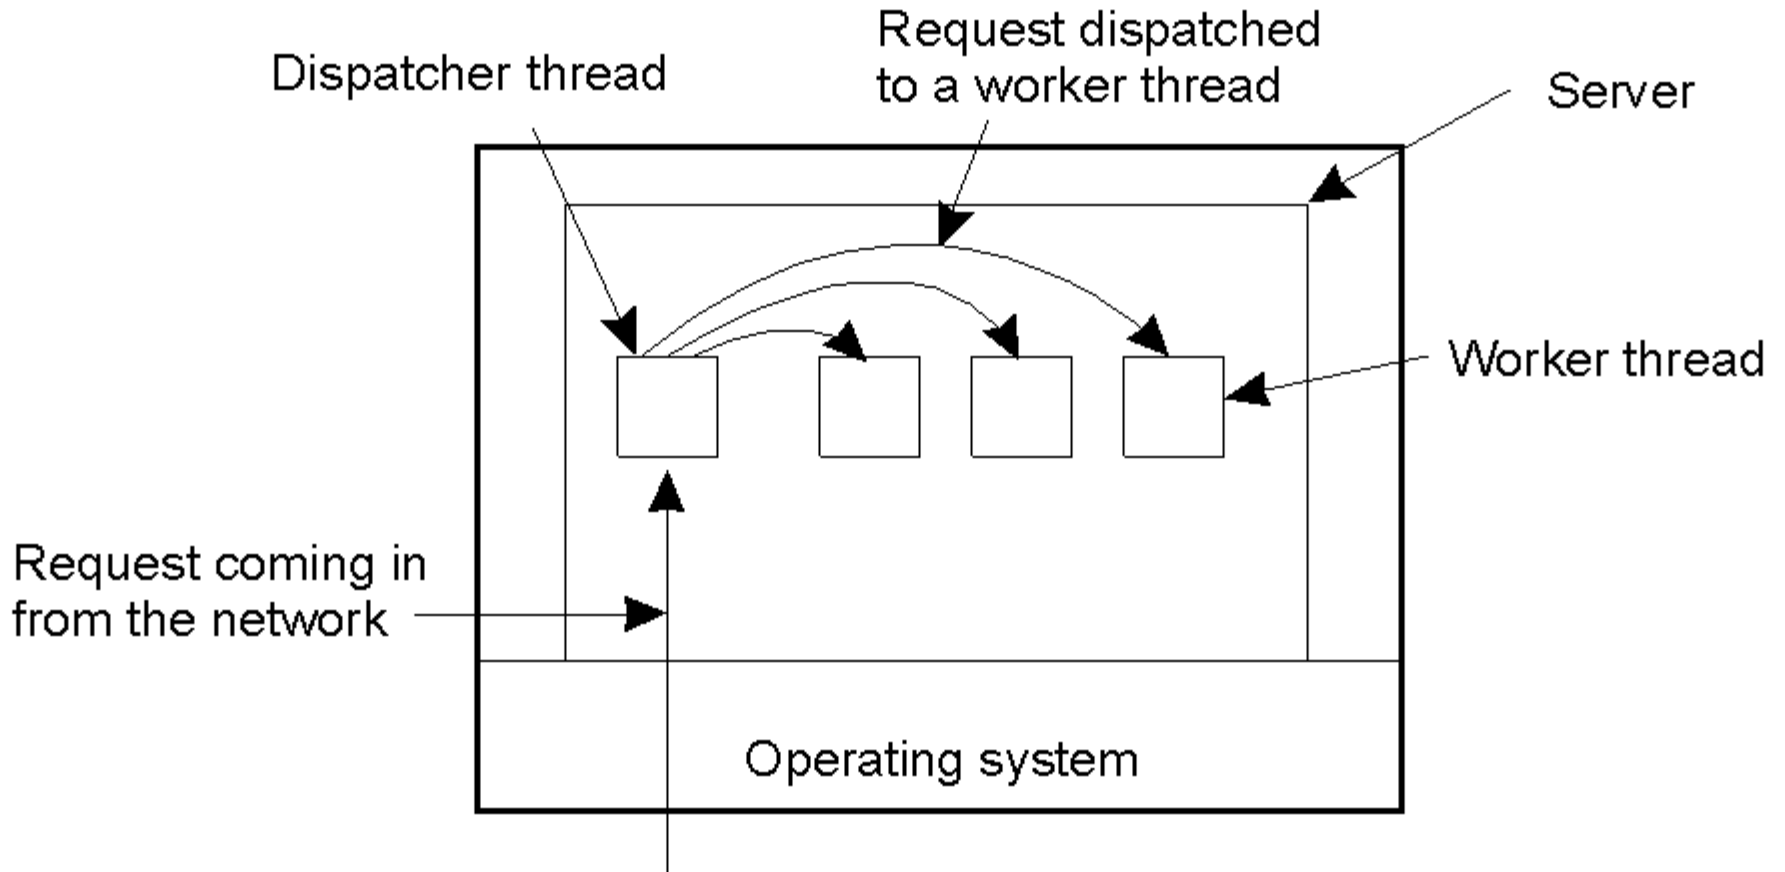
- A client does several calls at the same time, each one by a different thread.
- It then waits until all results have been returned.

Note: if calls are to different servers, we may have [linear speed-up](#).

# *Threads in Distributed Systems*

- Improve performance
  - Starting a thread is much cheaper than starting a new process.
  - Having a single-threaded server prohibits simple scale-up to a multiprocessor system.
  - As with clients: hide network latency by reacting to next request while previous one is being replied.
- Better structure
  - Most servers have high I/O demands. Using simple, well-understood blocking calls simplifies the overall structure.
  - Multithreaded programs tend to be smaller and easier to understand due to simplified flow of control.

# Multithreaded Servers (1)



A multithreaded server organized in a dispatcher/worker model.

## Multithreaded Servers (2)

<b>Model</b>	<b>Characteristics</b>
Threads	Parallelism, blocking system calls
Single-threaded process	No parallelism, blocking system calls
Finite-state machine	Parallelism, nonblocking system calls

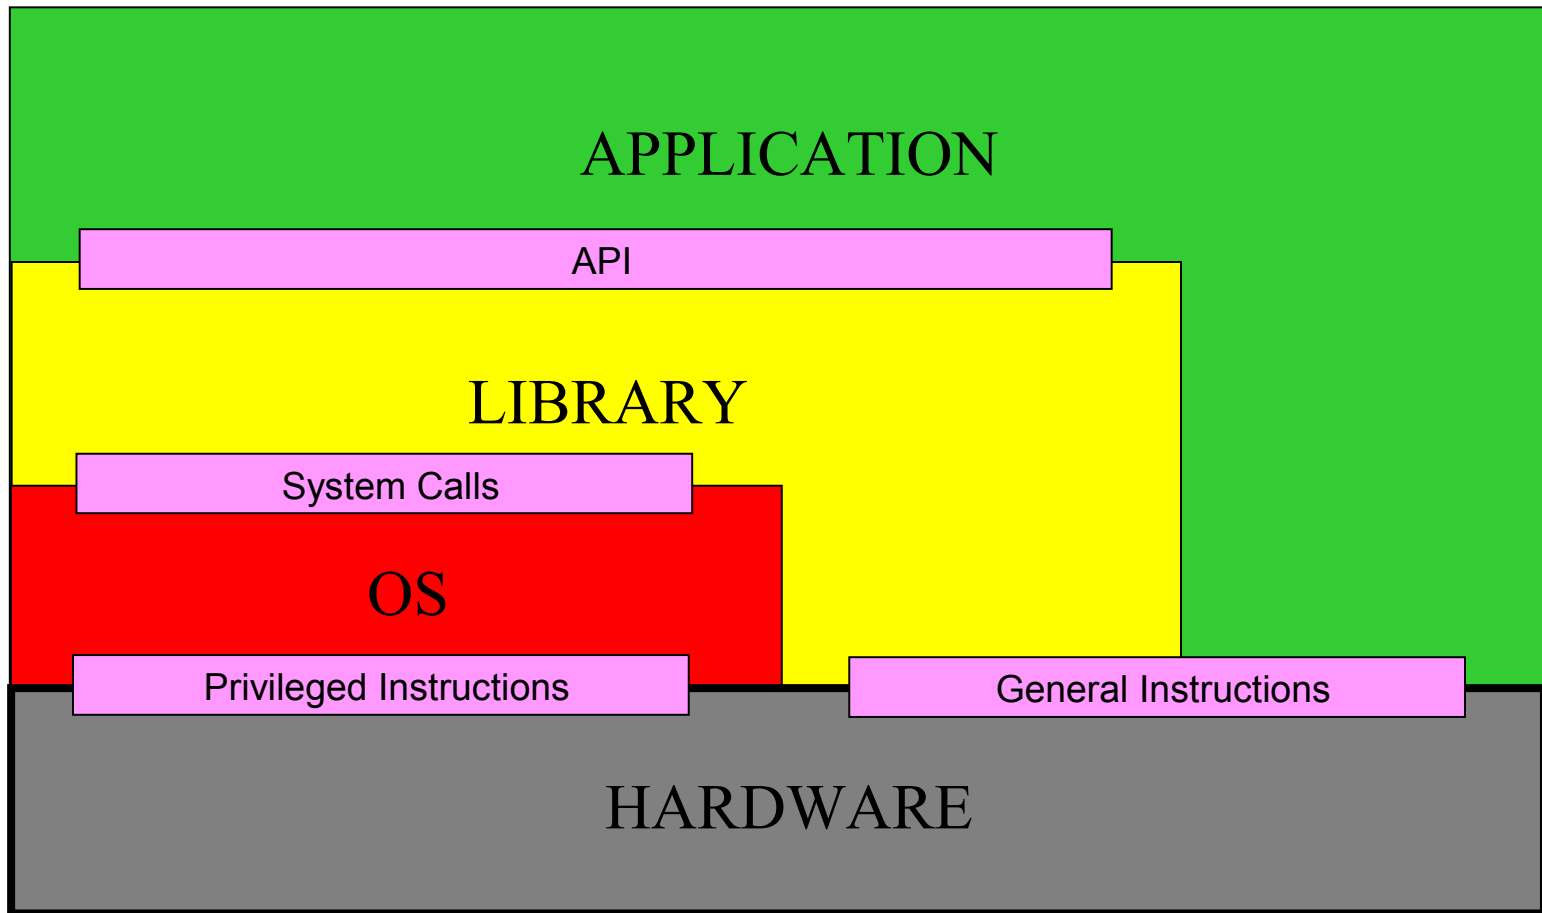
Three ways to construct a server.

# *Virtualization*

- Virtualization is becoming increasingly important:
  - Hardware changes faster than software
  - Ease of portability and code migration
  - Isolation of failing or attacked components



# Architecture of Virtual Machines



# *Types of Virtual Machines*


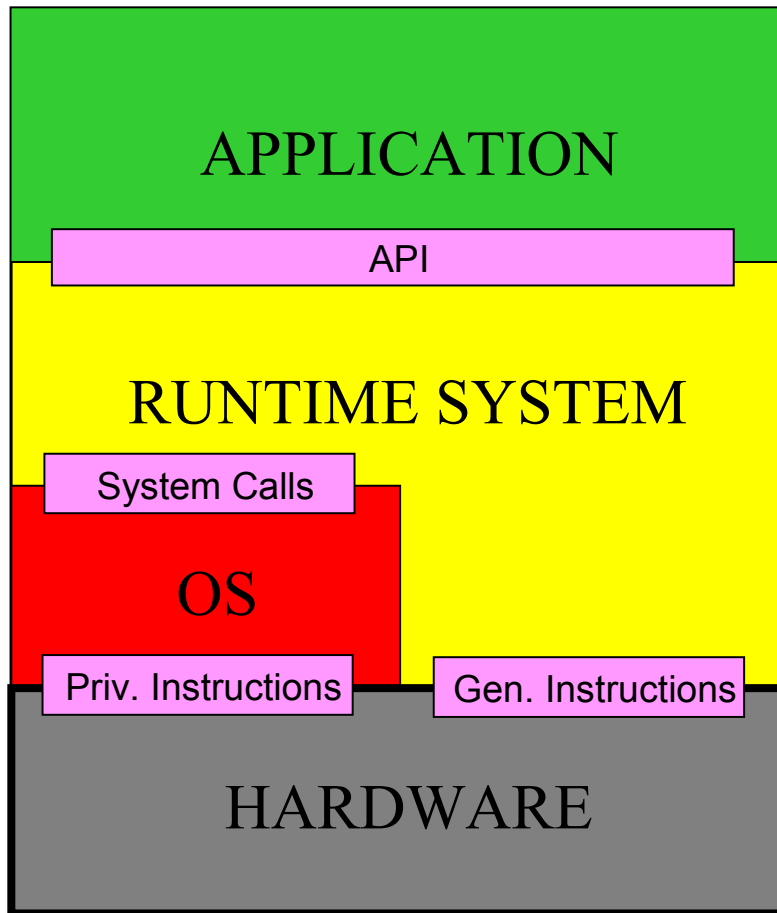
## Process Virtual Machine

- One VM per process
- Runtime system
- Interpreted or emulated instructions

## Virtual Machine Monitor

- One VM for more processes
- Layer that completely encapsulates the original h/w
- Interface to a virtual h/w

# Process VMs vs. VM Monitors

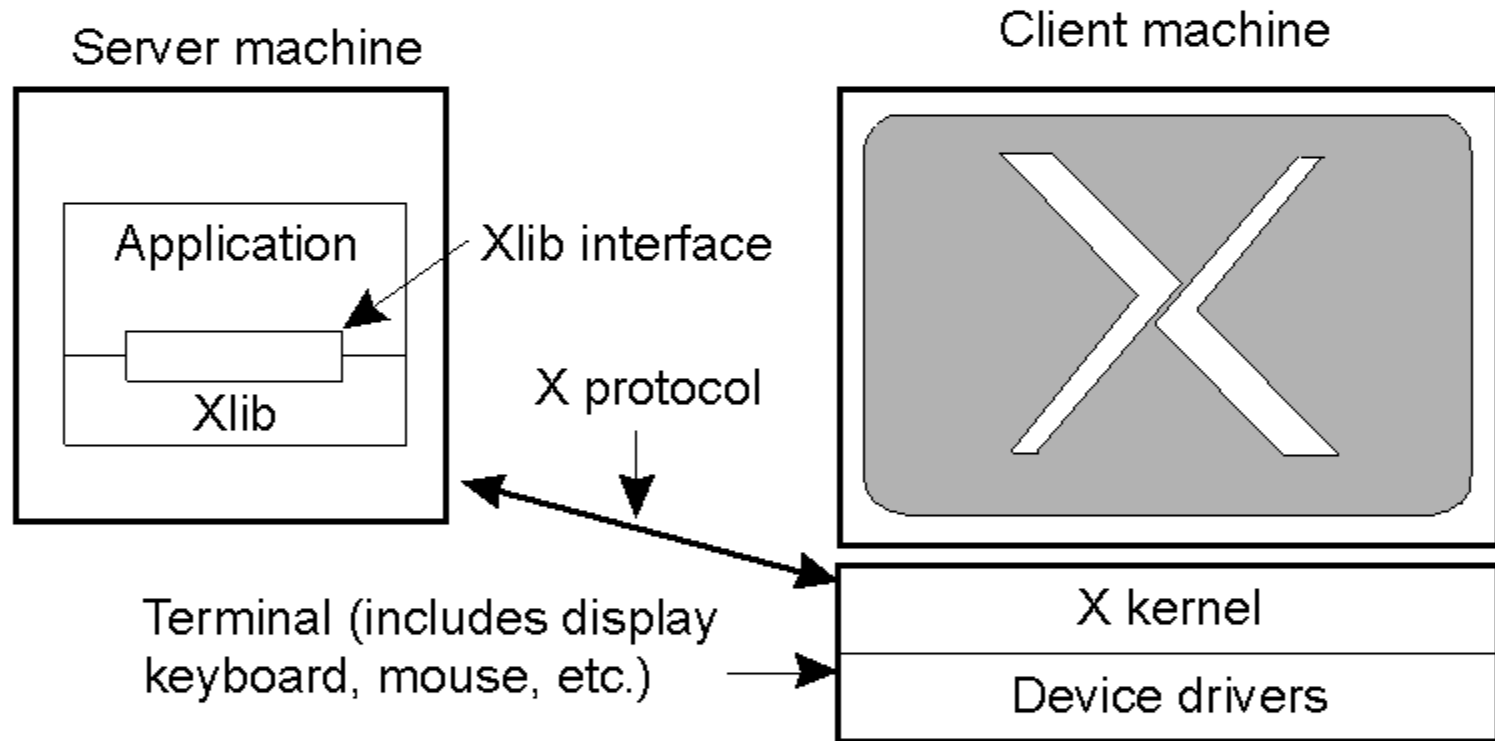


# *VM Monitors on Operating Systems*

We're seeing VMMs run on top of existing operating systems.

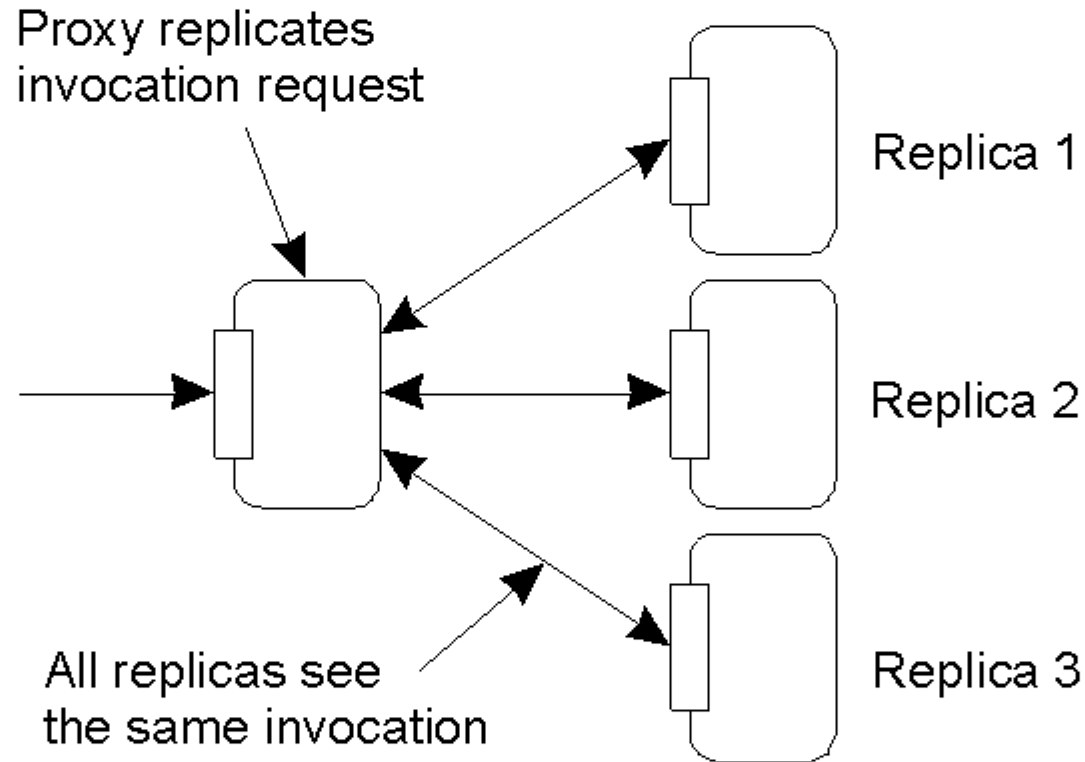
- Perform binary translation: while executing an application on an operating system, translate instructions to that of the underlying machine.
- Distinguish sensitive instructions: traps to the original kernel (think of system calls, or privileged instructions).
- Sensitive instructions are replaced with calls to the VMM.

# Clients: User Interfaces



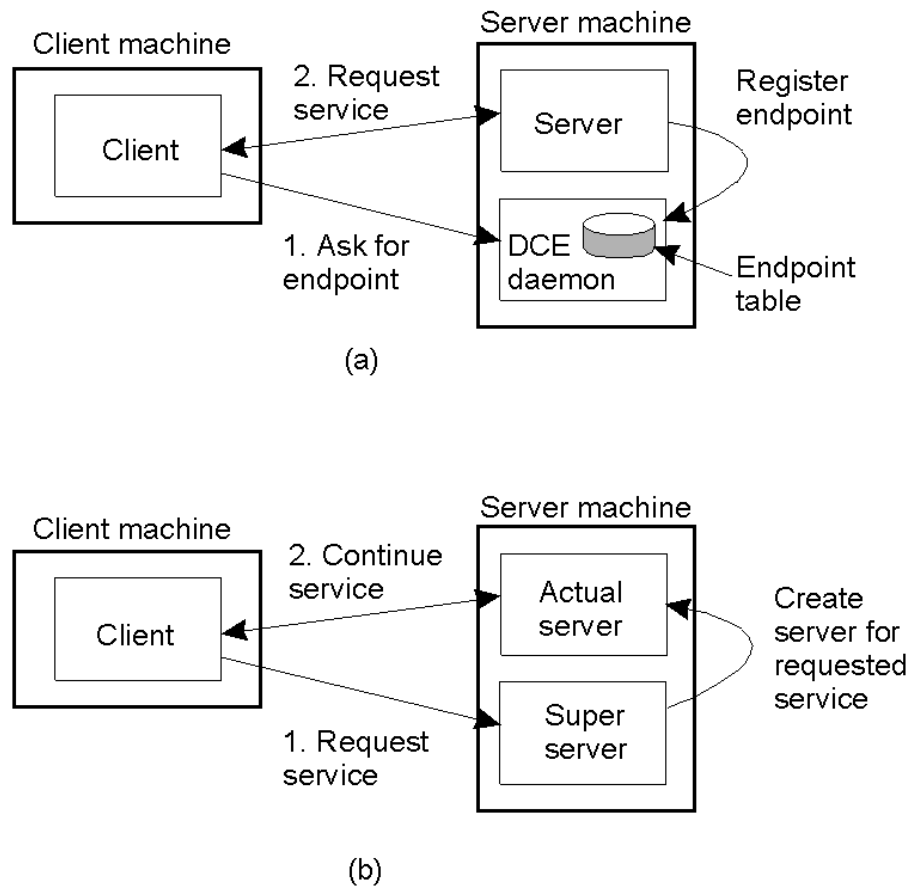
The basic organization of the X Window System

# Client-Side Software for Distribution Transparency



A possible approach to transparent replication of a remote object using a client-side solution.

# Servers: General Design Issues



- a) Client-to-server binding using a daemon as in DCE
- b) Client-to-server binding using a superserver as in UNIX

# *Out-of-Band Communication*

**Issue:** Is it possible to interrupt a server once it has accepted (or is in the process of accepting) a service request?

- Solution 1: Use a separate port for urgent data:
  - Server has a separate thread/process for urgent messages
  - Urgent message comes in  $\Rightarrow$  associated request put on hold
  - Note: we require OS supports priority-based scheduling
- Solution 2: Use out-of-band communication facilities of the transport layer:
  - Example: TCP allows for urgent messages in same connection
  - Urgent messages can be caught using OS signaling techniques



# *Servers and State*

## Stateless servers

- Never keep accurate information about the status of a client after having handled a request:
- Don't record whether a file has been opened (simply close it again after access)
- Don't promise to invalidate a client's cache
- Don't keep track of your clients

## Consequences

- Clients and servers are completely independent
- State inconsistencies due to client or server crashes are reduced
- Possible loss of performance because, e.g., a server cannot anticipate client behavior (think of prefetching file blocks)

# *Servers and State*

Stateful servers: Keep track of the status of their clients:

- Record that a file has been opened, so that prefetching can be done
- Know which data a client has cached, and allow clients to keep local copies of shared data

Observation

- The performance of stateful servers can be extremely high, provided clients are allowed to keep local copies. As it turns out, reliability is not a major problem.

*Thank you for your attention*

