

Figure 18.1 Mach supports operating systems, databases and other subsystems

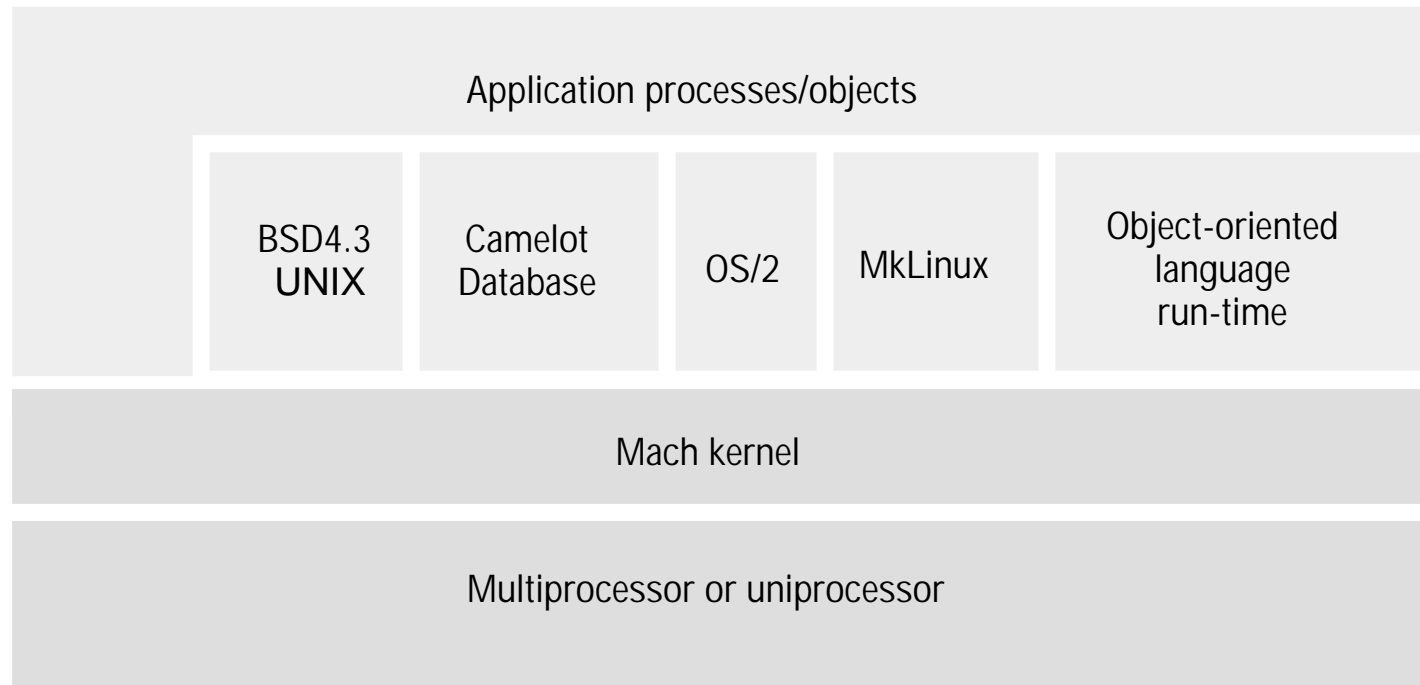


Figure 18.2 Mach tasks, threads and communication

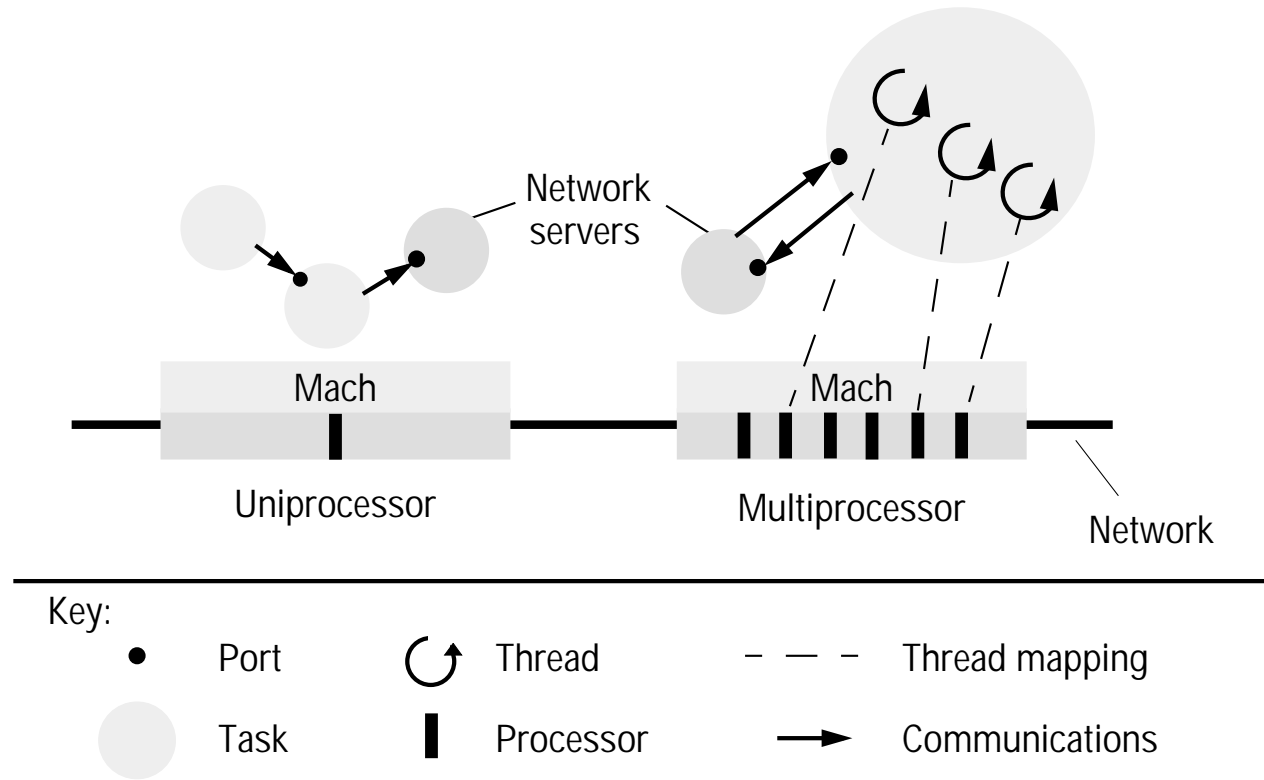


Figure 18.3 A task's port name space

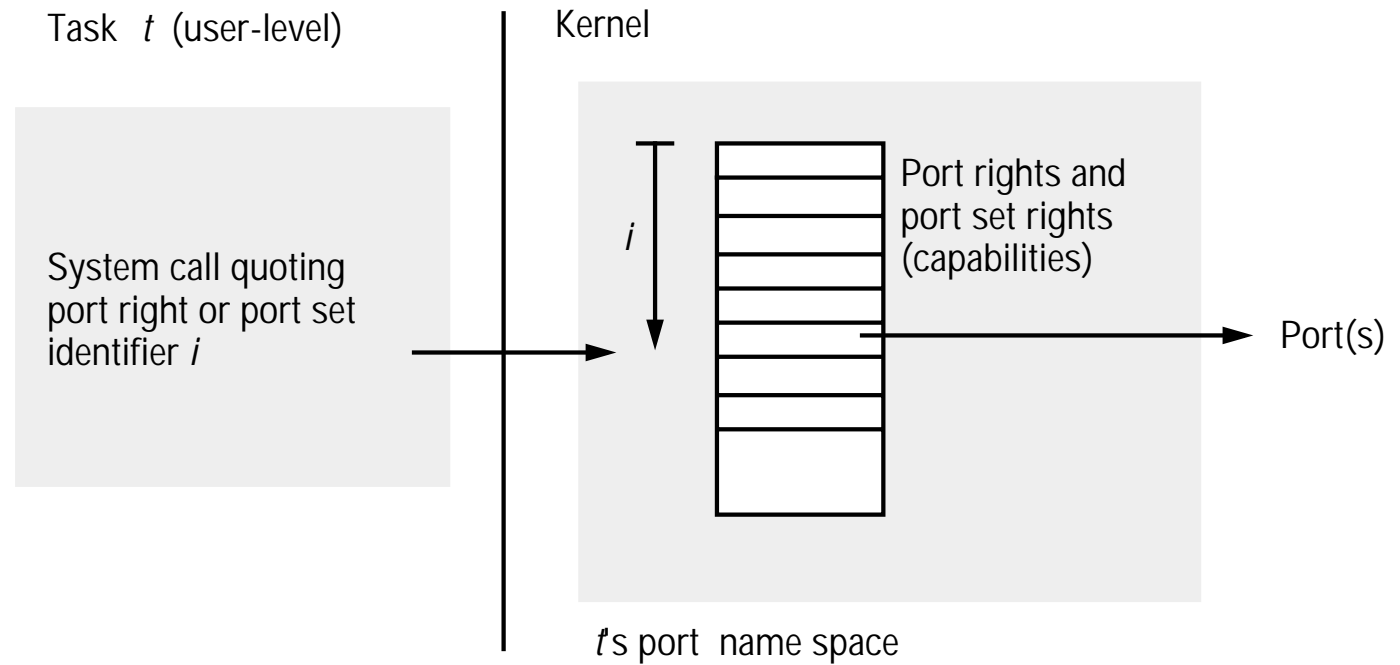


Figure 18.4 Task and thread creation

task_create(parent_task, inherit_memory, child_task)

parent_task is the task used as a blueprint in the creation of the new task, *inherit_memory* specifies whether the child should inherit the address space of its parent or be assigned an empty address space, *child_task* is the identifier of the new task.

thread_create(parent_task, child_thread)

parent_task is the task in which the new thread is to be created, *child_thread* is the identifier of the new thread. The new thread has no execution state and is suspended.

thread_set_state(thread, flavour, new_state, count)

thread is the thread to be supplied with execution state, *flavour* specifies the machine architecture, *new_state* specifies the state (such as the program counter and stack pointer), *count* is the size of the state.

thread_resume(thread)

This is used to resume the suspended thread identified by *thread*.

Figure 18.5 A Mach message containing port rights and out-of-line data

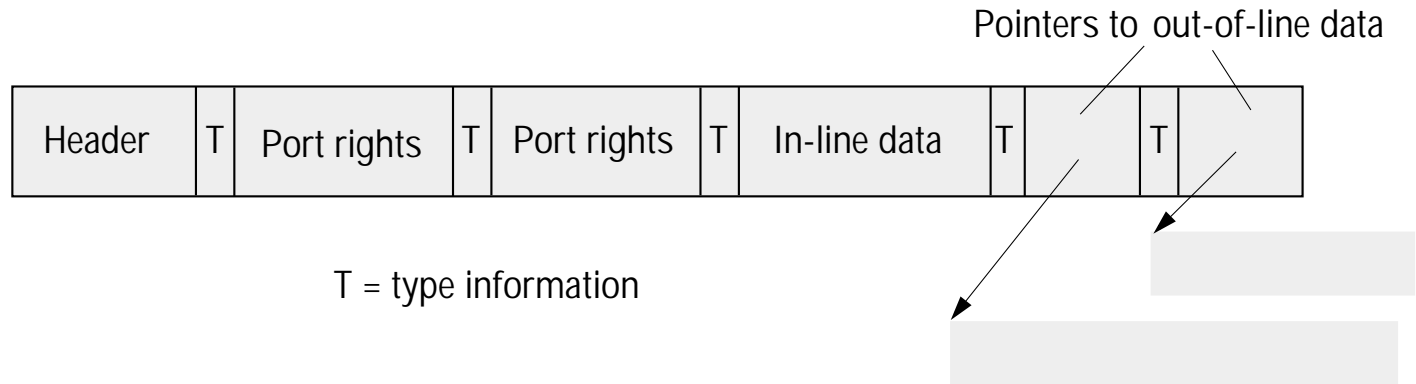


Figure 18.6 Network communication in Mach

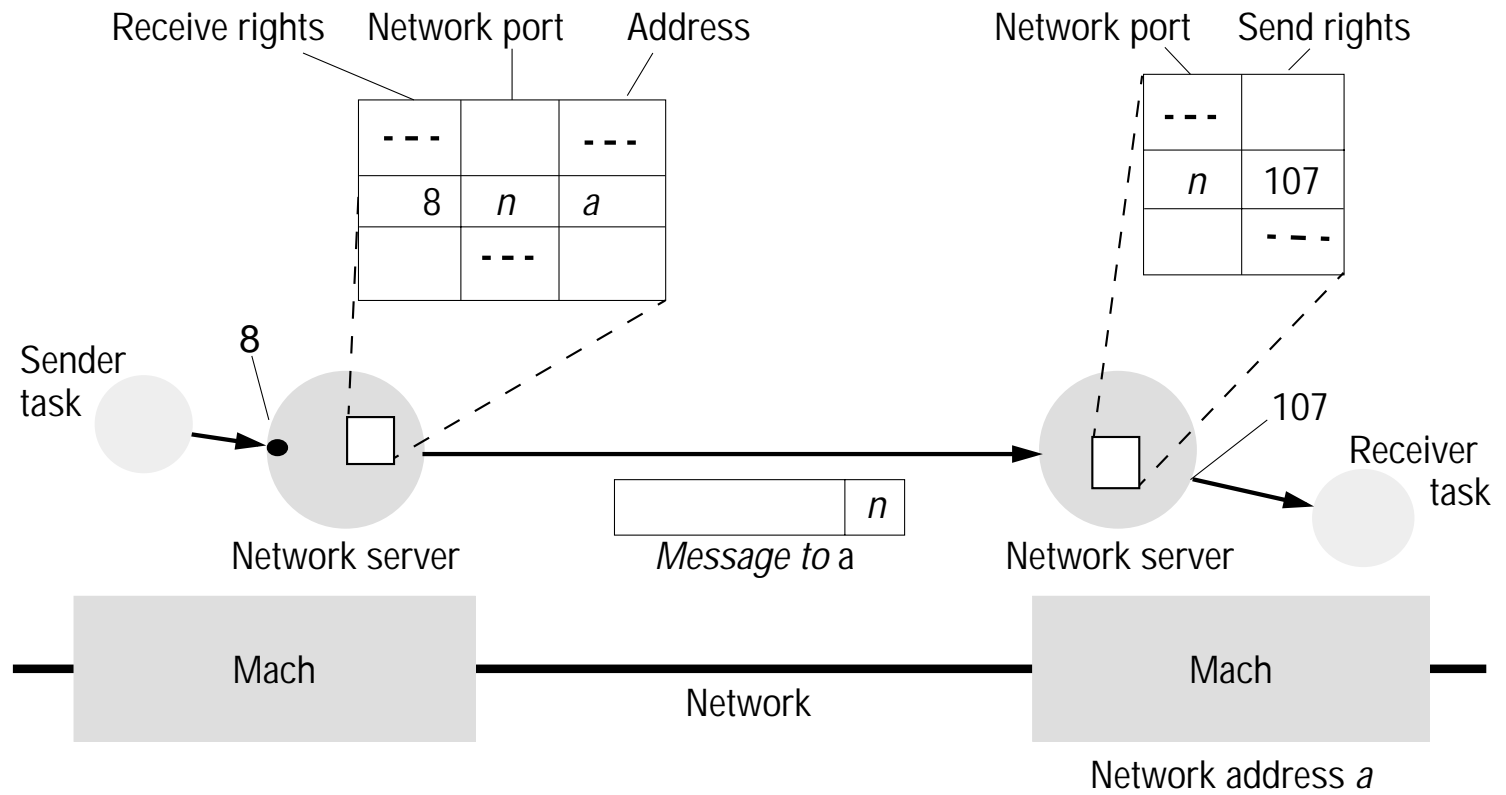


Figure 18.7 Copy-on-write overheads

<i>Region size</i>	<i>Simple copy</i>	<i>Create region</i>	<i>Amount of data copied (on writing)</i>		
			<i>0 kilobytes (0 pages)</i>	<i>8 kilobytes (1 page)</i>	<i>256 kilobytes (32 pages)</i>
8 kilobytes	1.4	1.57	2.7	4.82	–
256 kilobytes	44.8	1.81	2.9	5.12	66.4

Note: all times are in milliseconds.

Figure 18.8 External pager




Figure 18.9 External pager messages

<i>Event</i>	<i>Sender</i>	<i>Message</i>
	K → EP	<i>memory_object_init</i>
<i>vm_map</i> called by task	EP → K	<i>memory_object_set_attributes</i> , or
	EP → K	<i>memory_object_data_error</i>
Task page-faults when no data frame exists	K → EP	<i>memory_object_data_request</i>
	EP → K	<i>memory_object_data_provided</i> , or
	EP → K	<i>memory_object_data_unavailable</i>
Kernel writes modified page to persistent store	K → EP	<i>memory_object_data_write</i>
External pager directs kernel to write page/set access permissions	EP → K	<i>memory_object_lock_request</i>
	K → EP	<i>memory_object_lock_completed</i>
Task page-faults when insufficient page access	K → EP	<i>memory_object_data_unlock</i>
	EP → K	<i>memory_object_lock_request</i>
Memory object no longer mapped	K → EP	<i>memory_object_terminate</i>
External pager withdraws memory object	EP → K	<i>memory_object_destroy</i>
	K → EP	<i>memory_object_terminate</i>