

Resource Containers

Gaurav Banga

Peter Druschel

Jeffrey C. Mogul

Presented by Michael Roitzsch



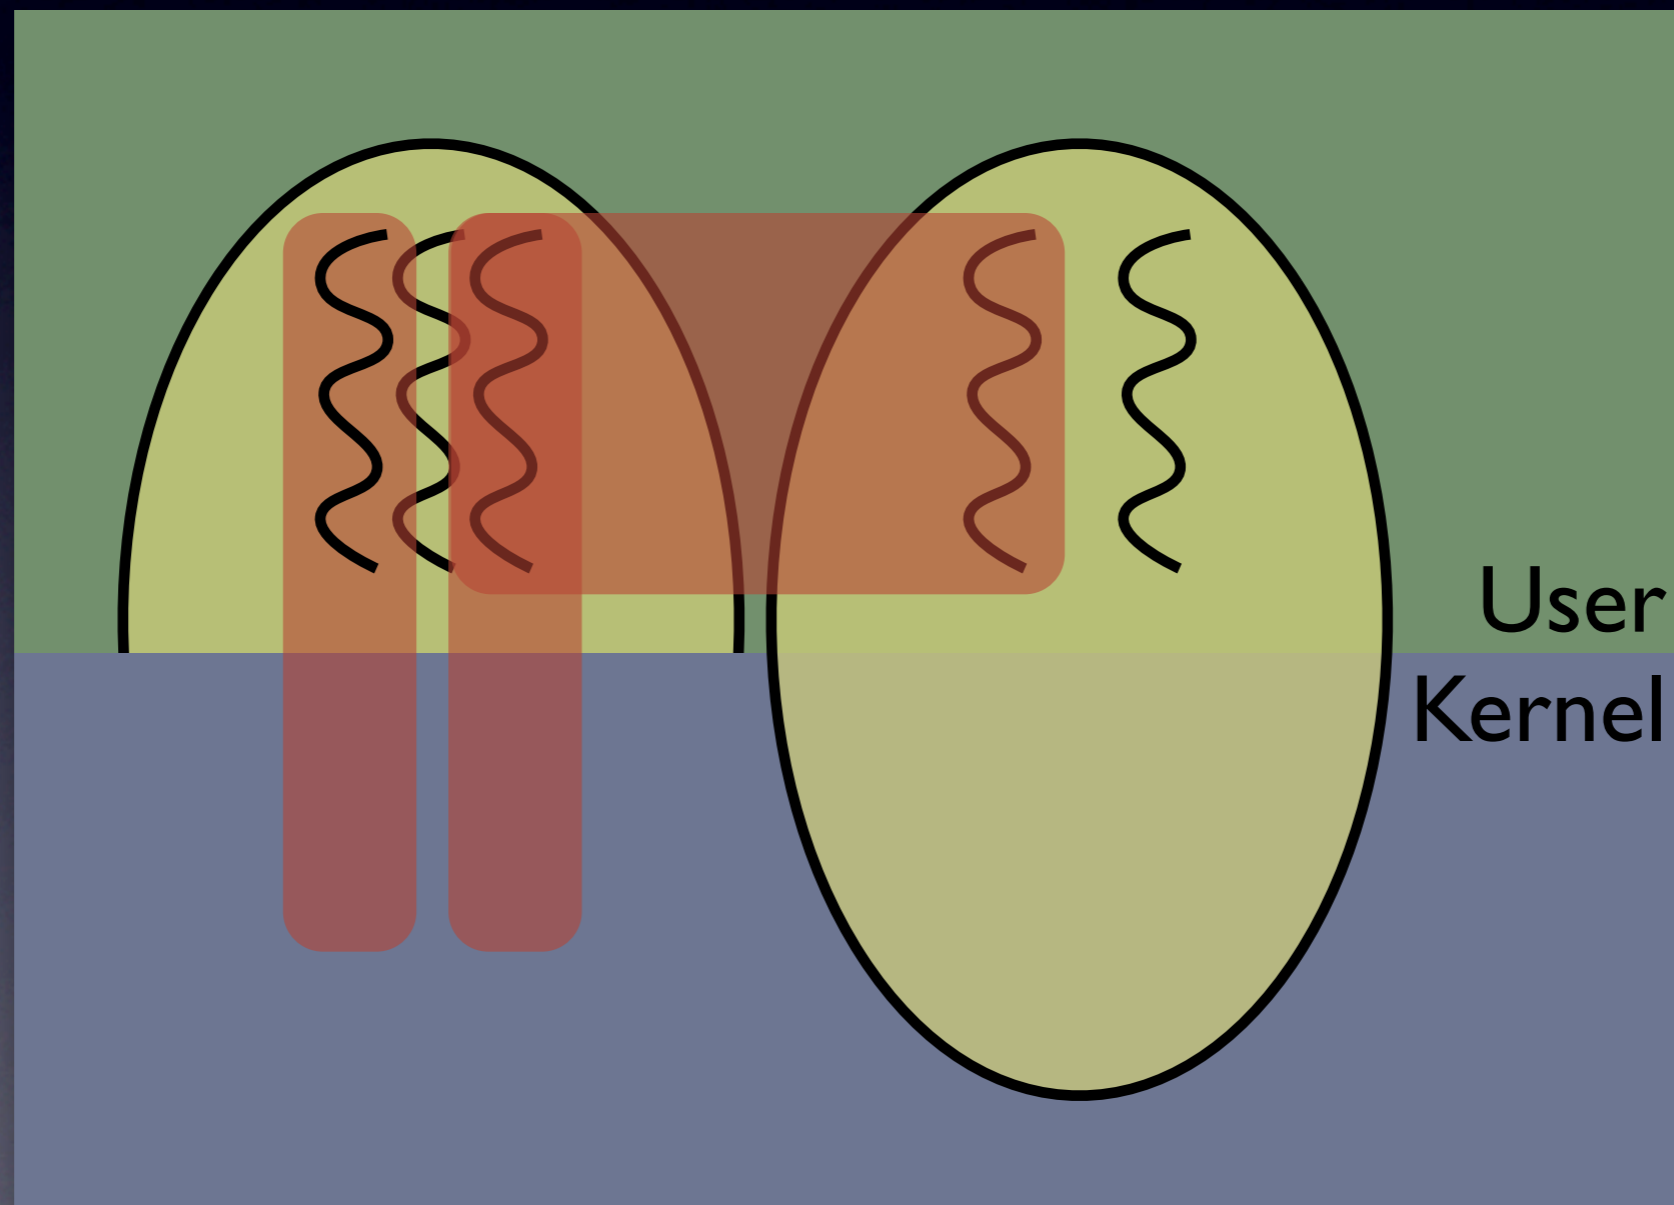
Observations

- mismatch: OS'es resource management design assumptions and behavior of modern server applications
- no control over resource consumption by the kernel on behalf of the application
- difficulty to express priority policies, QoS
- denial of service

Web Server Example

HTTP Connection Handling	CGI Handling
process forking	CGI process forking
pre-forked processes	persistent CGI processes
event driven	library based
multi-threaded	

Current Resource Management



Dual Use of Processes




Resource Containers

- abstract operating system entity
- contains resources used for an activity,
kernel accounts against resource containers
- arbitrary relationships between protection domains, threads, resource containers
- mechanism to be used with resource management policy
- form a hierarchy

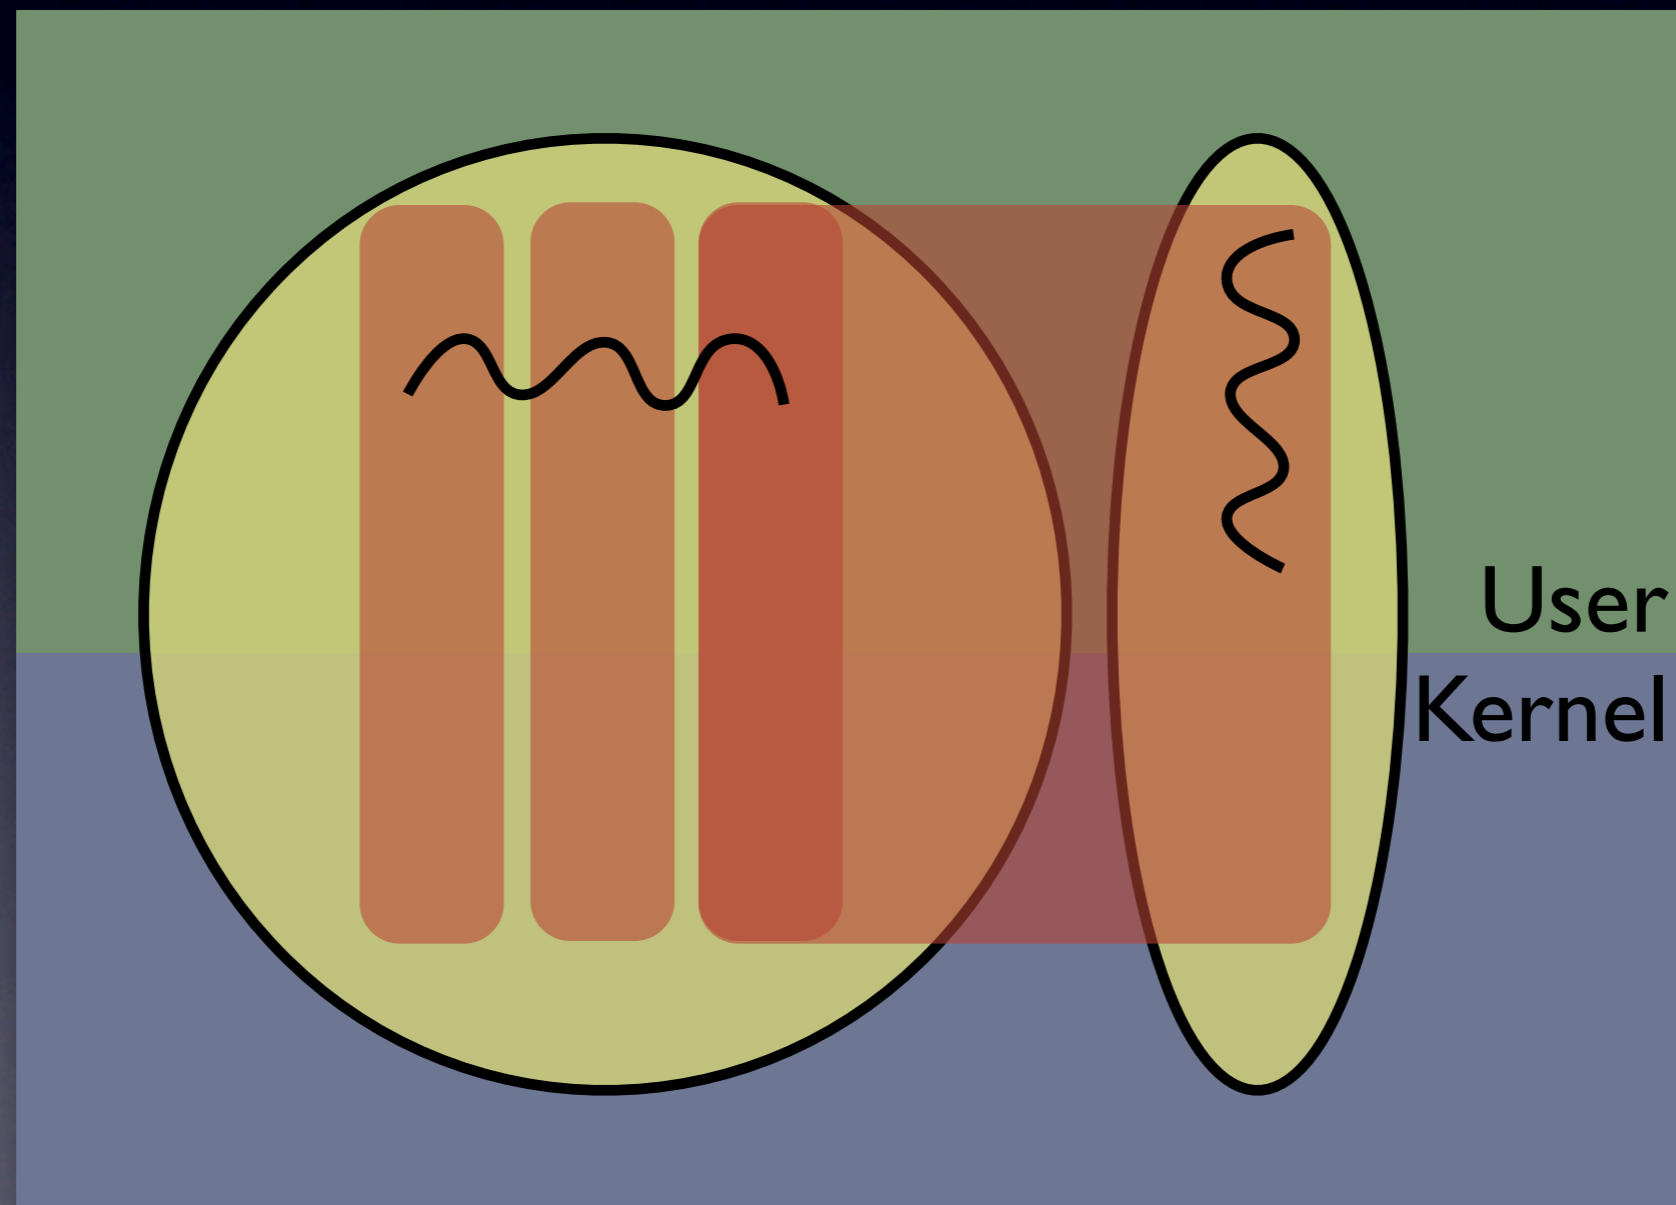
Operations


- creating a new container
- set a container's parent
- container release
- sharing containers between processes
- container attributes
- container usage information
- binding a thread to a container
- reset the scheduler binding
- binding a socket or a file to a container

Multi-Threaded Server



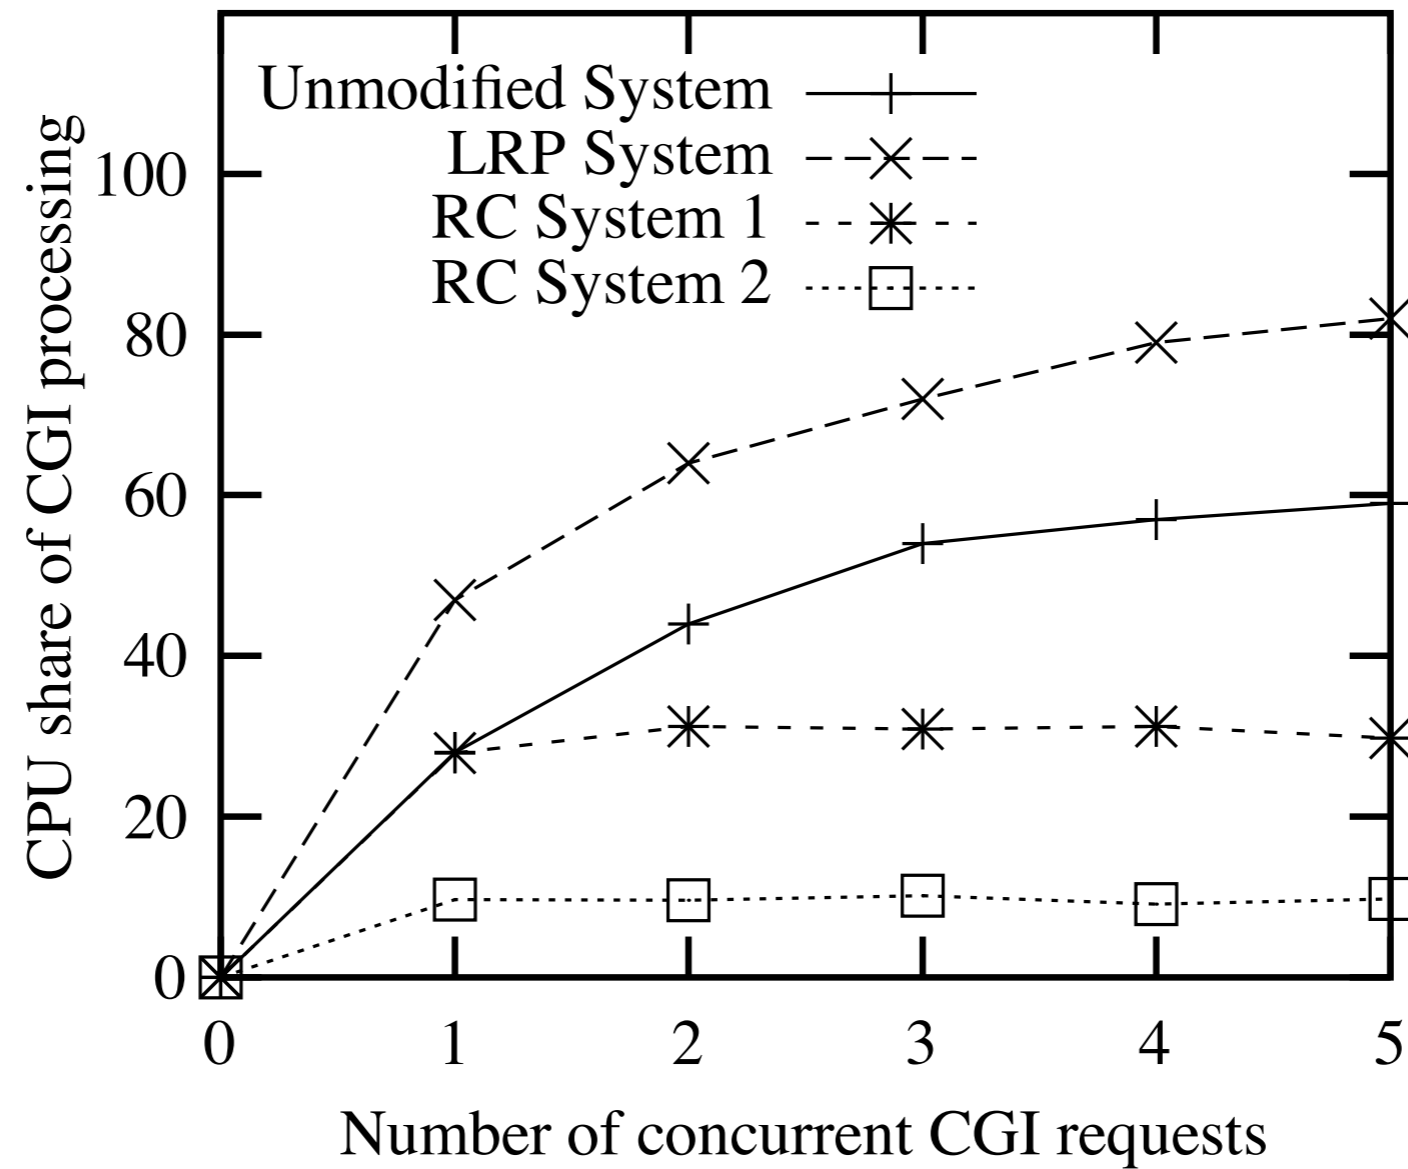
Event-Driven Server







Number of concurrent low-priority clients

0 2 10 12 30 32 30 32







Discussion

- How do microkernels solve such problems?
- Why can a container be reparented or even made parent-less?
- How to account kernel resource consumption properly?
- Is the motivation for the problem still valid?