

Orran Krieger, Marc Auslander, Bryan Rosenberg,
Robert Wisniewski, Jimi Xenidis, Dilma Da Silva,
Michal Ostrowski, Jonathan Appavoo, Maria Butrico,
Mark Mergen, Amos Waterland, Volkmar Uhlig

K42: Building a Complete Operating System

Presented by Thomas Aschauer
at the Software Systems Seminar,
Department of Computer Sciences,
University of Salzburg



Contents

- History & Goals
- Key Concepts
- Conclusion

K42 History & Goals



K42 History: Technical Predictions in 1996

- Windows dominant
- Multiprocessors more important
- Increasing OS maintenance costs
- Customizability & extensibility critical
- All machines 64 bit in five years



Research Goals

- Performance & scalability
 - Small & large multiprocessor
- Customizability
- Applicability
- Wide availability
 - Open source & maintainability

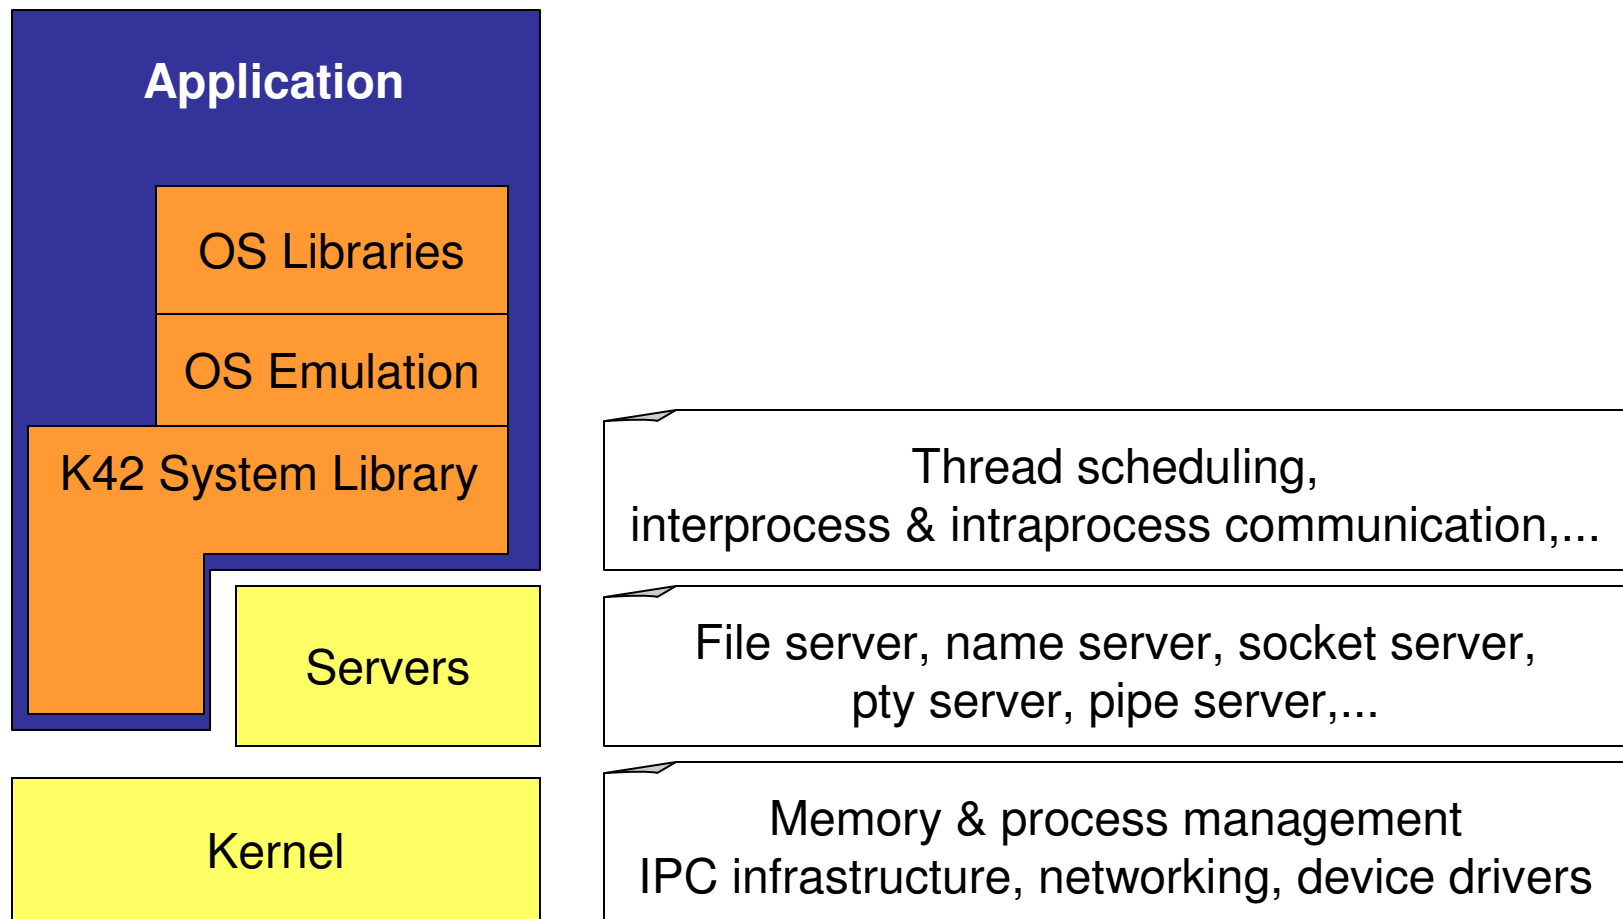


Technical Directions

- Start from scratch
- Exokernel design
- User level implementations
- OO Design

K42 Key Concepts

System Structure






Key concept (a): Clustered Object

- Object in OO-sense
- Resides on one or more processors
- Services requests for one or more processors

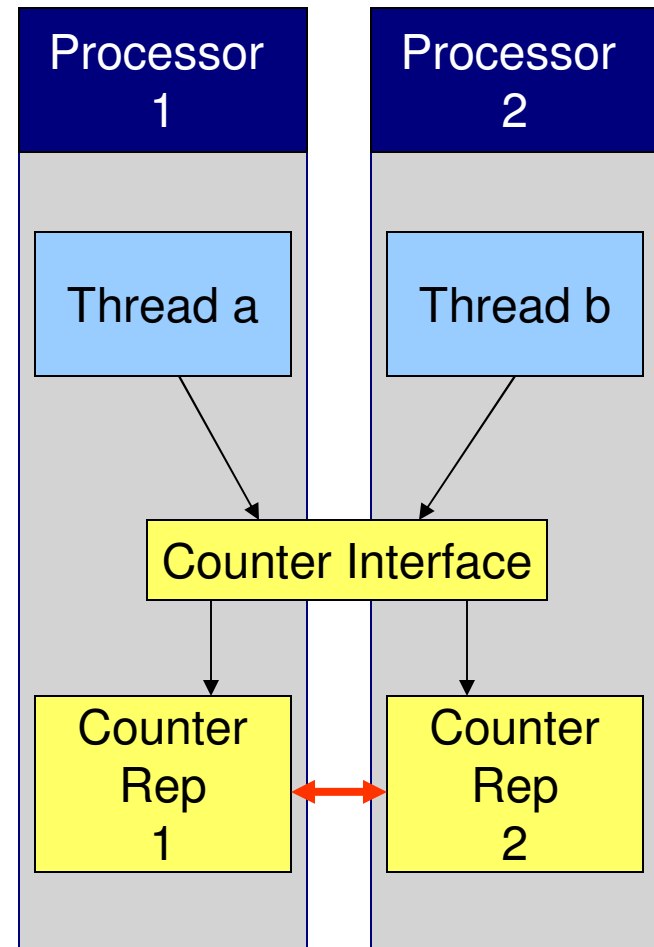
Example: Clustered Object “Counter”

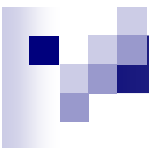
■ Interface



■ Implementation

- (a) Shared value
- (b) Distributed value





Key Concept (b): Customization

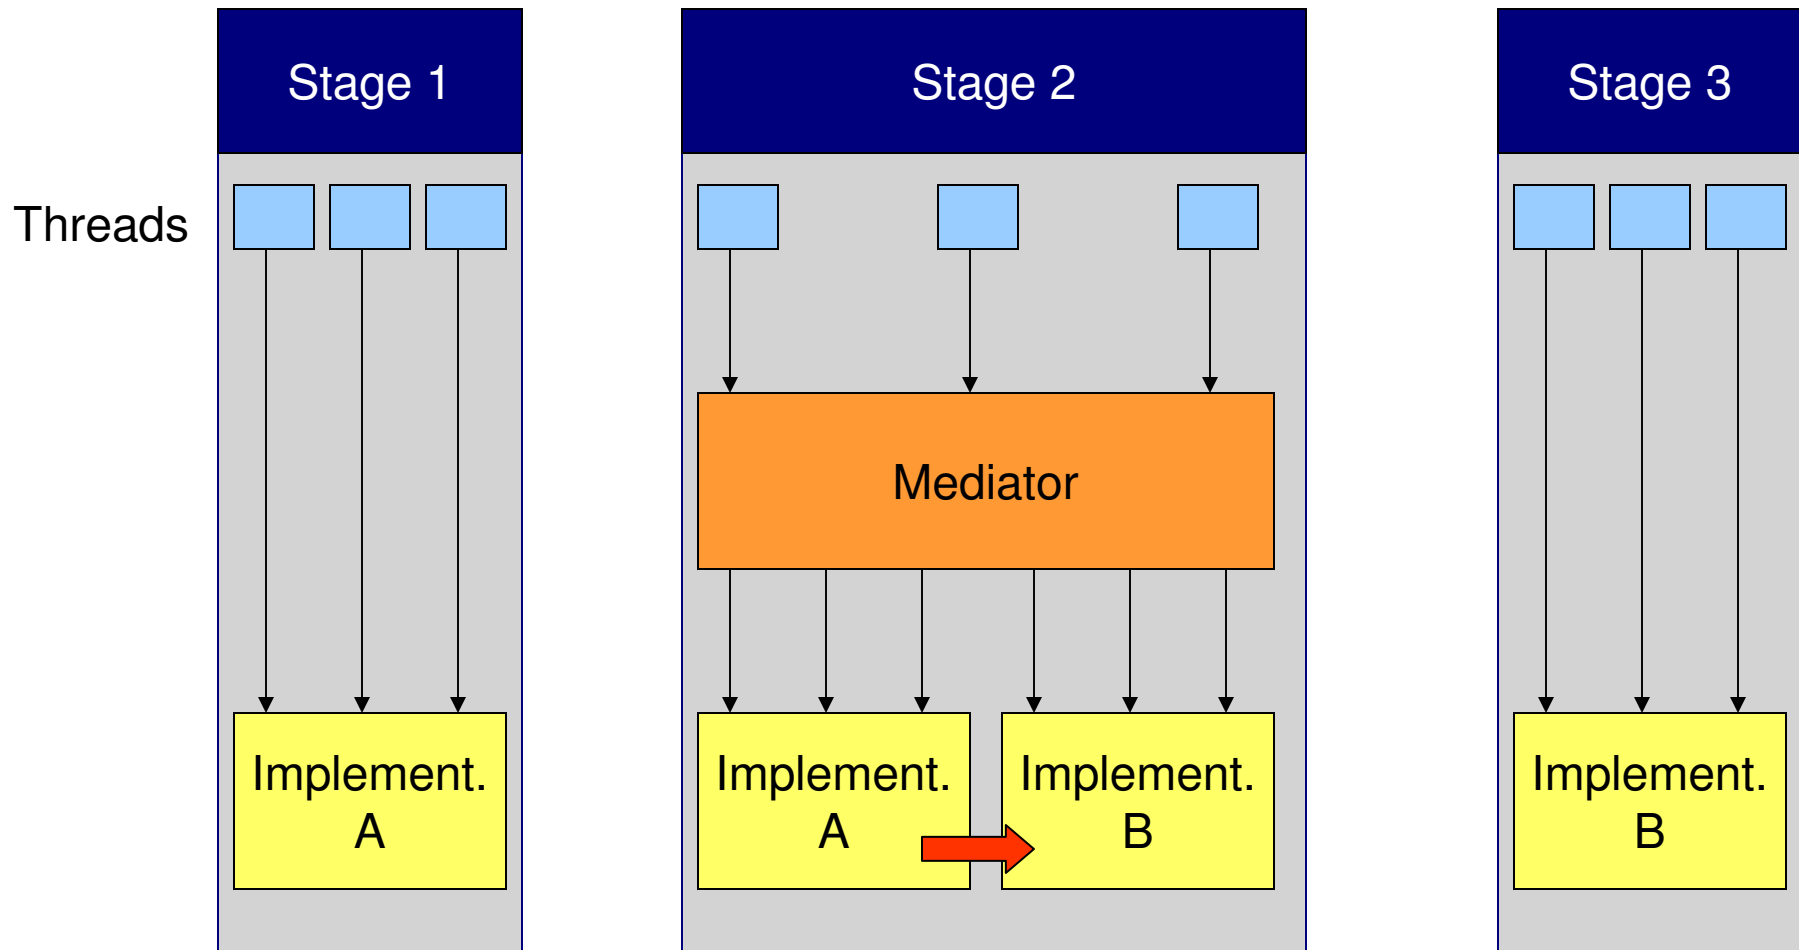
- Hot swapping

- Replace active object with new implementation

- Dynamic upgrade

- Replace all objects providing a certain service
- Uses hot swapping

Example: Hot Swapping

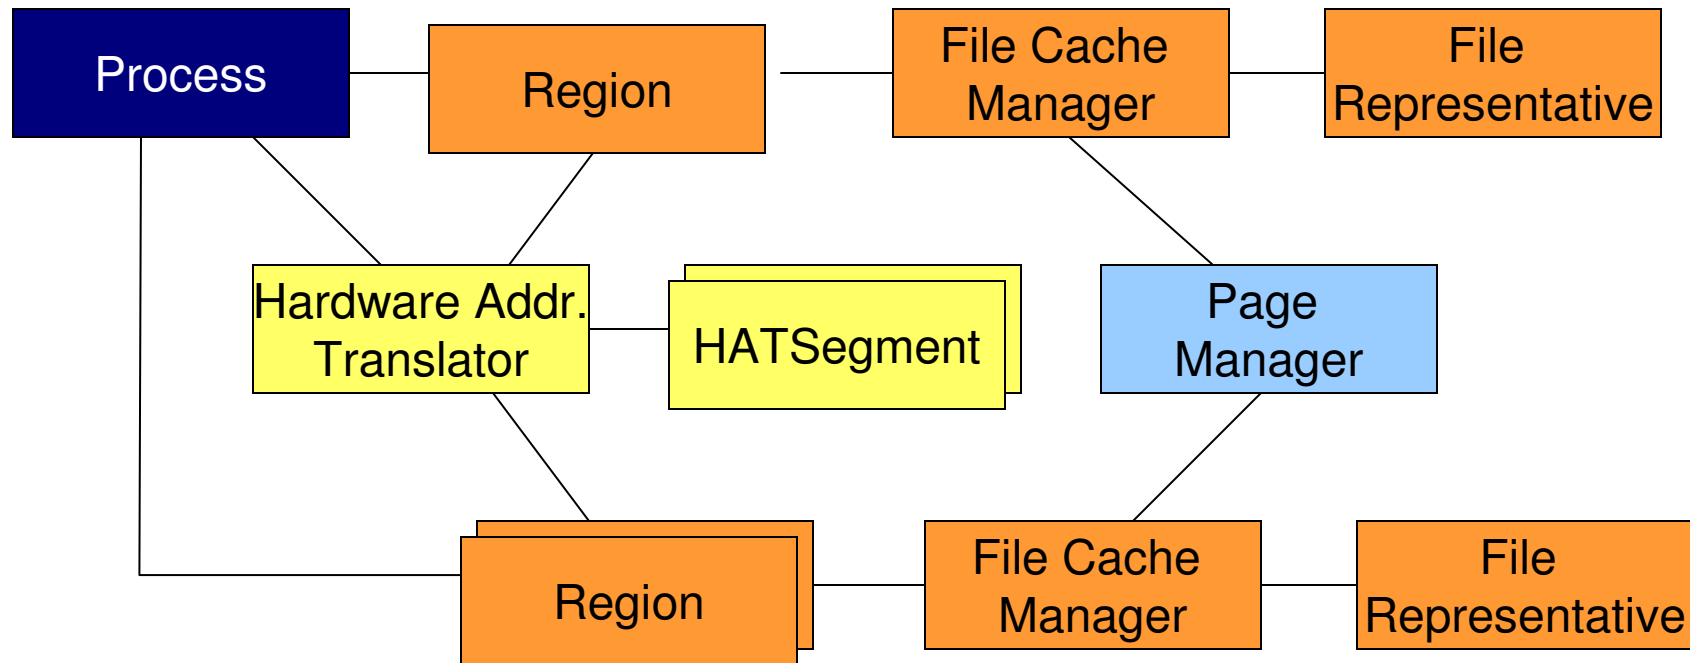




Key Concept (c): OO Design

- Use OO design
 - Whenever applicable
 - “One instance per resource”
- Avoid global data structures and policies
 - They do not scale well
- Separate service mechanism from policy
 - Can be customized independently

Example: OO Memory Management

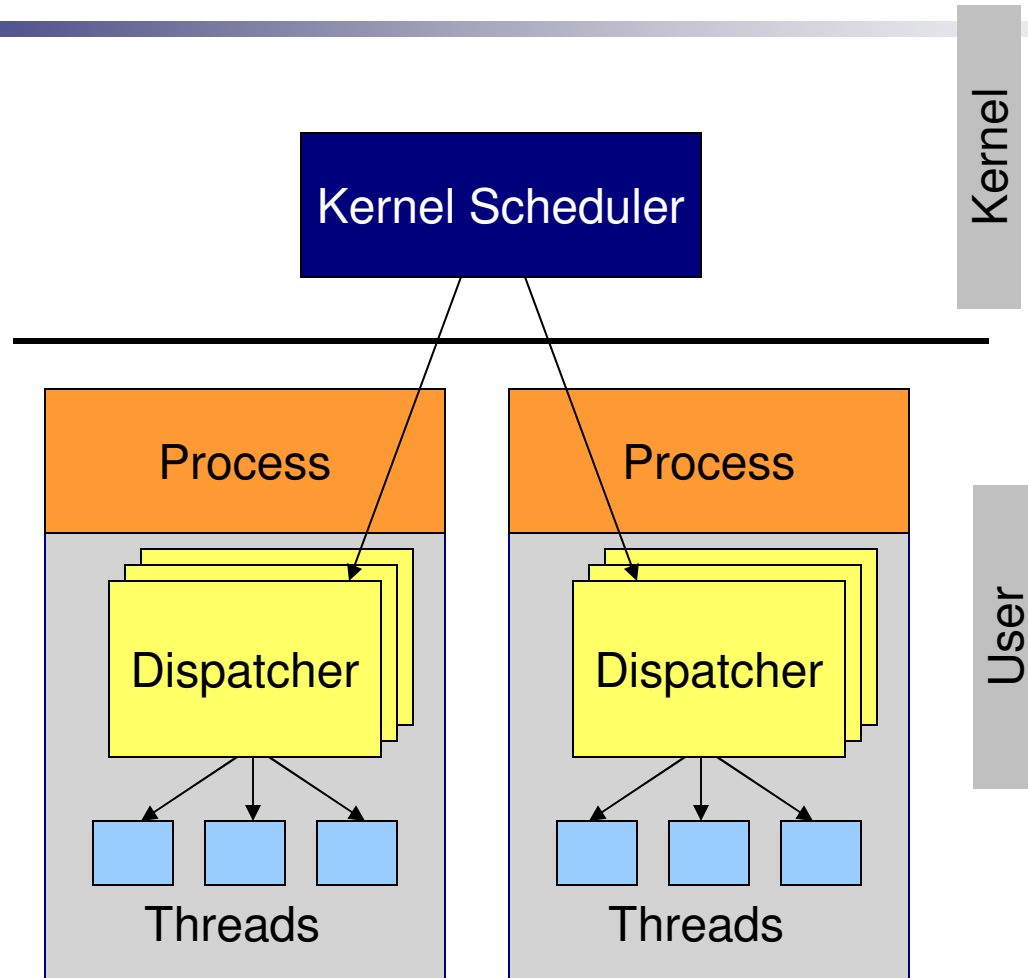




Key Concept (d): User-level Implementation of Kernel Functionality

- Goals
 - Avoid system call
 - Minimize kernel resources
- E.g. Scheduling

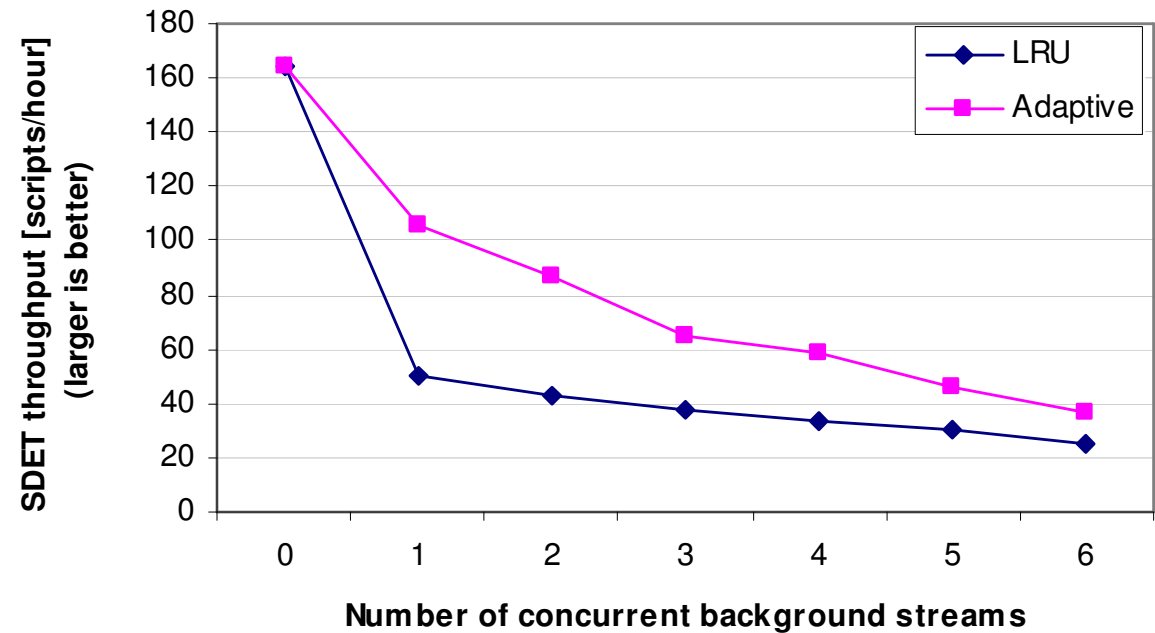
Example: User-Level Scheduling



Concluding Remarks

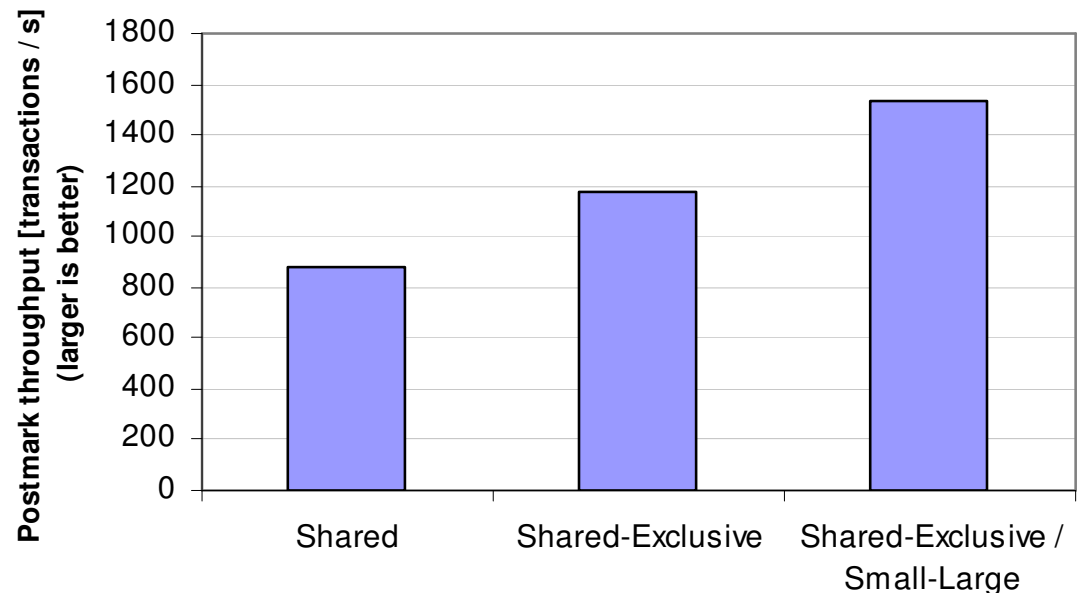
Hot-Swapping Performance Gain [1]

- Adaptive page replacement
- Streaming background applications
- Monitor page usage pattern
- Hot swap to sequential optimized FCM




Hot-Swapping Performance Gain [2]

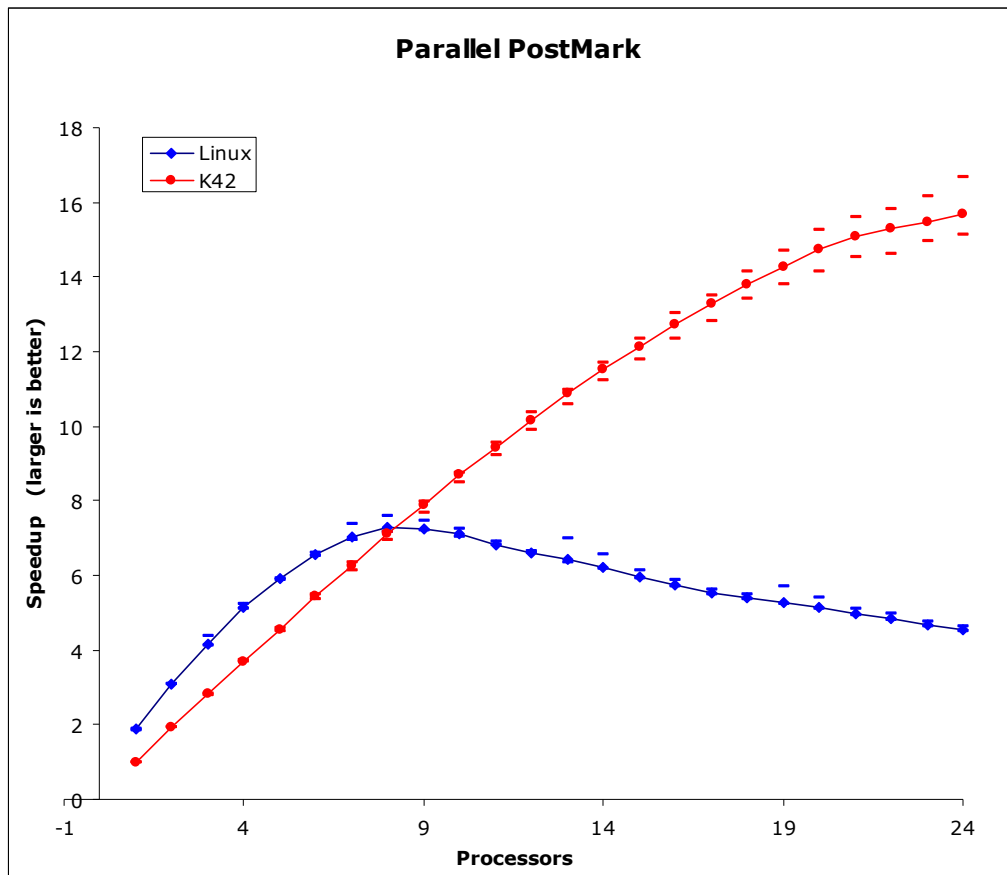
- Adaptive file cache
 - 1: Default implementation
 - 2: Optimized for non-shared files
 - 3: Cache small files in application address space




K42 vs. Linux 2.4.19 [1]



K42 vs. Linux 2.4.19 [2]



K42 vs. Linux 2.4.19 [3]





Achieved Goals

- Scalability & customizability
 - Through OO design
- Customizability
 - Hot swapping & dynamic upgrade
- Applicability
 - Linux compatibility
- Wide availability
 - Open source
- Still open
 - Maintainability



Thank You!



References

- <http://www.research.ibm.com/K42>
- Oran Krieger, Marc Auslander, Bryan Rosenberg, Robert W. Wisniewski, Jimi Xenidis, Dilma Da Silva, Michal Ostrowski, Jonathan Appavoo, Maria Butrico, Mark Mergen, Amos Waterland, Volkmar Uhlig: K42: Building a Complete Operating System. Proceedings of the 2006 EuroSys Conference, Leuven, Belgium.
- Orran Krieger, Marc Auslander, Bryan Rosenberg, Robert Wisniewski, Jimi Xenidis, Dilma Da Silva, Michal Ostrowski, Jonathan Appavoo, Maria Butrico, Mark Mergen, Amos Waterland, Volkmar Uhlig: K42: Building a Complete Operating System. Presentation at the 2006 EuroSys Conference, Leuven, Belgium. See <http://www.cs.kuleuven.ac.be/conference/EuroSys2006/>

Backup Slides



Discussion on OO

■ Advantages

- Easy to provide special implementations
- Per-instance resource management enables autonomic system optimization

■ Disadvantages

- Scattered control flow
- Performance overhead
 - Outperformed
- Really more maintainable?



Linux Compatibility

- Linux API & ABI compatibility
- glibc
 - Unmodified version
 - System trap reflection
 - Modified version
 - Directly calls K42 system lib in user space
- Kernel support
 - Directly linking TCP/IP stack, file systems, drivers, ...
 - Provides Linux code environment
 - Significant maintenance effort needed