



IEEE
SecDev | 2020



Refactoring the FreeBSD Kernel with Checked C


Junhan Duan, Yudi Yang, Jie Zhou, and John Criswell



 #IEEESecDev

 <https://secdev.ieee.org/2020>

C Lacks Memory Safety!



What if:

1. $a[4] \leftarrow 10$
2. $a[5] \leftarrow$ malicious function

Buffer Overflow Attack

Checked C: Spatial Memory Safety for C

- Compatible with existing C code
- Designed for incremental change

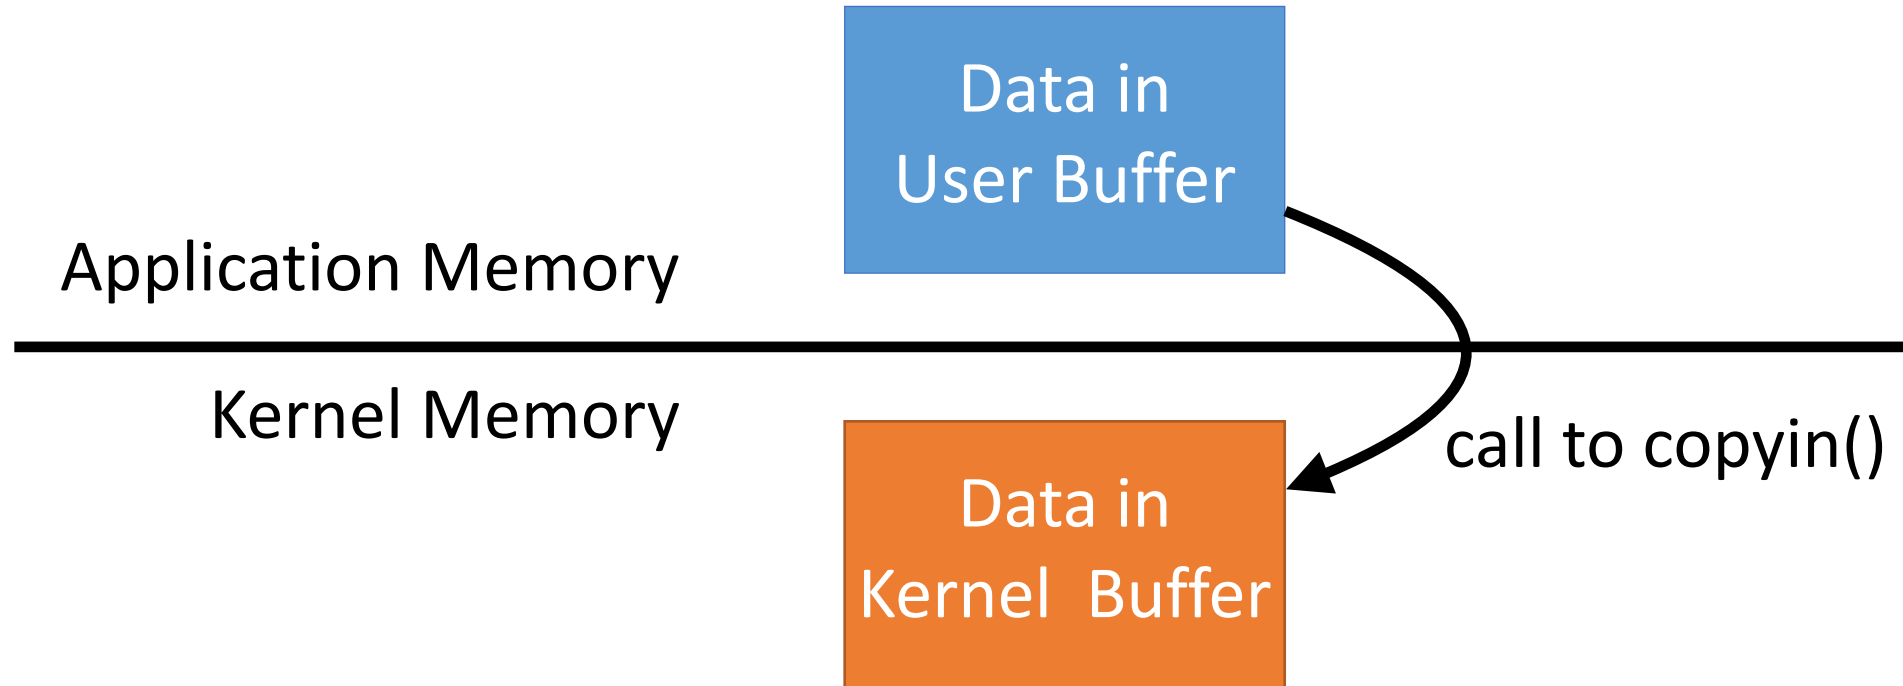
Checked C: Spatial Memory Safety for C

- New pointer types: `ptr<T>` and `array_ptr<T>`
- Bounds Safe Interface
 - Allow checked and unchecked code to inter-operate
 - Polymorphic pointer types for function arguments
- Checked Region
 - All code within region guaranteed to be spatially safe
 - Code within are restricted to use checked pointers only

Our Refactoring Effort

- Kernel Support Routine: Memory manipulations
 - memcpy(), memmove(), memset()
 - malloc(), free()
- System Calls
 - Calls to copyin(), copyout(), copyinstr()
- TCP/IP Stack
 - mbuf code, IP and UDP packet manipulation code

Copying Data Between Applications and Kernel



TCP / IP Stack

- mbuf package
 - network packet manipulation
- ip_input() & ip_output() & relevant functions
 - process ip headers
- udp_input() & udp_output() & relevant functions
 - process udp headers

Reflections on Checked C

- **Benefits**

- Compatible
- Incremental

- **Limitations**

- Compiler: $p + x - x \neq p$
- Macro arguments cannot take both checked and unchecked pointers

Reflections on Checked C

Compatible & Incremental

- Supports C pointer types
- Bounds Safe Interface
 - Inter-operation between checked and unchecked code
- Checked Region
 - Ensure spatial memory safety

Reflections on Checked C

- **Limitations**

- Compiler: lacking support for inferring expression equivalence
 - $p + x - x \neq p$
- Macros cannot take both checked and unchecked pointers
- Casting cannot be used as function inputs

Inconvenience with Macros

Original

- `#define memcpy ...`

Modified

- `#define memcpy ...`
- `#define memcpy_checked ...`

Evaluation

- Refactoring Effort
- Performance Overhead
 - System Calls
 - UDP/IP
- Code Size Overhead

Evaluation - Refactoring Effort

Added	Deleted	Files Modified
1,779	587	61

Lines of code modified on FreeBSD kernel


Evaluation - Refactoring Effort

- Refactored by two undergraduate students
- Spent 3 months, 7.3 hours a week
- Around 3 hours a week spent on implementation for the project

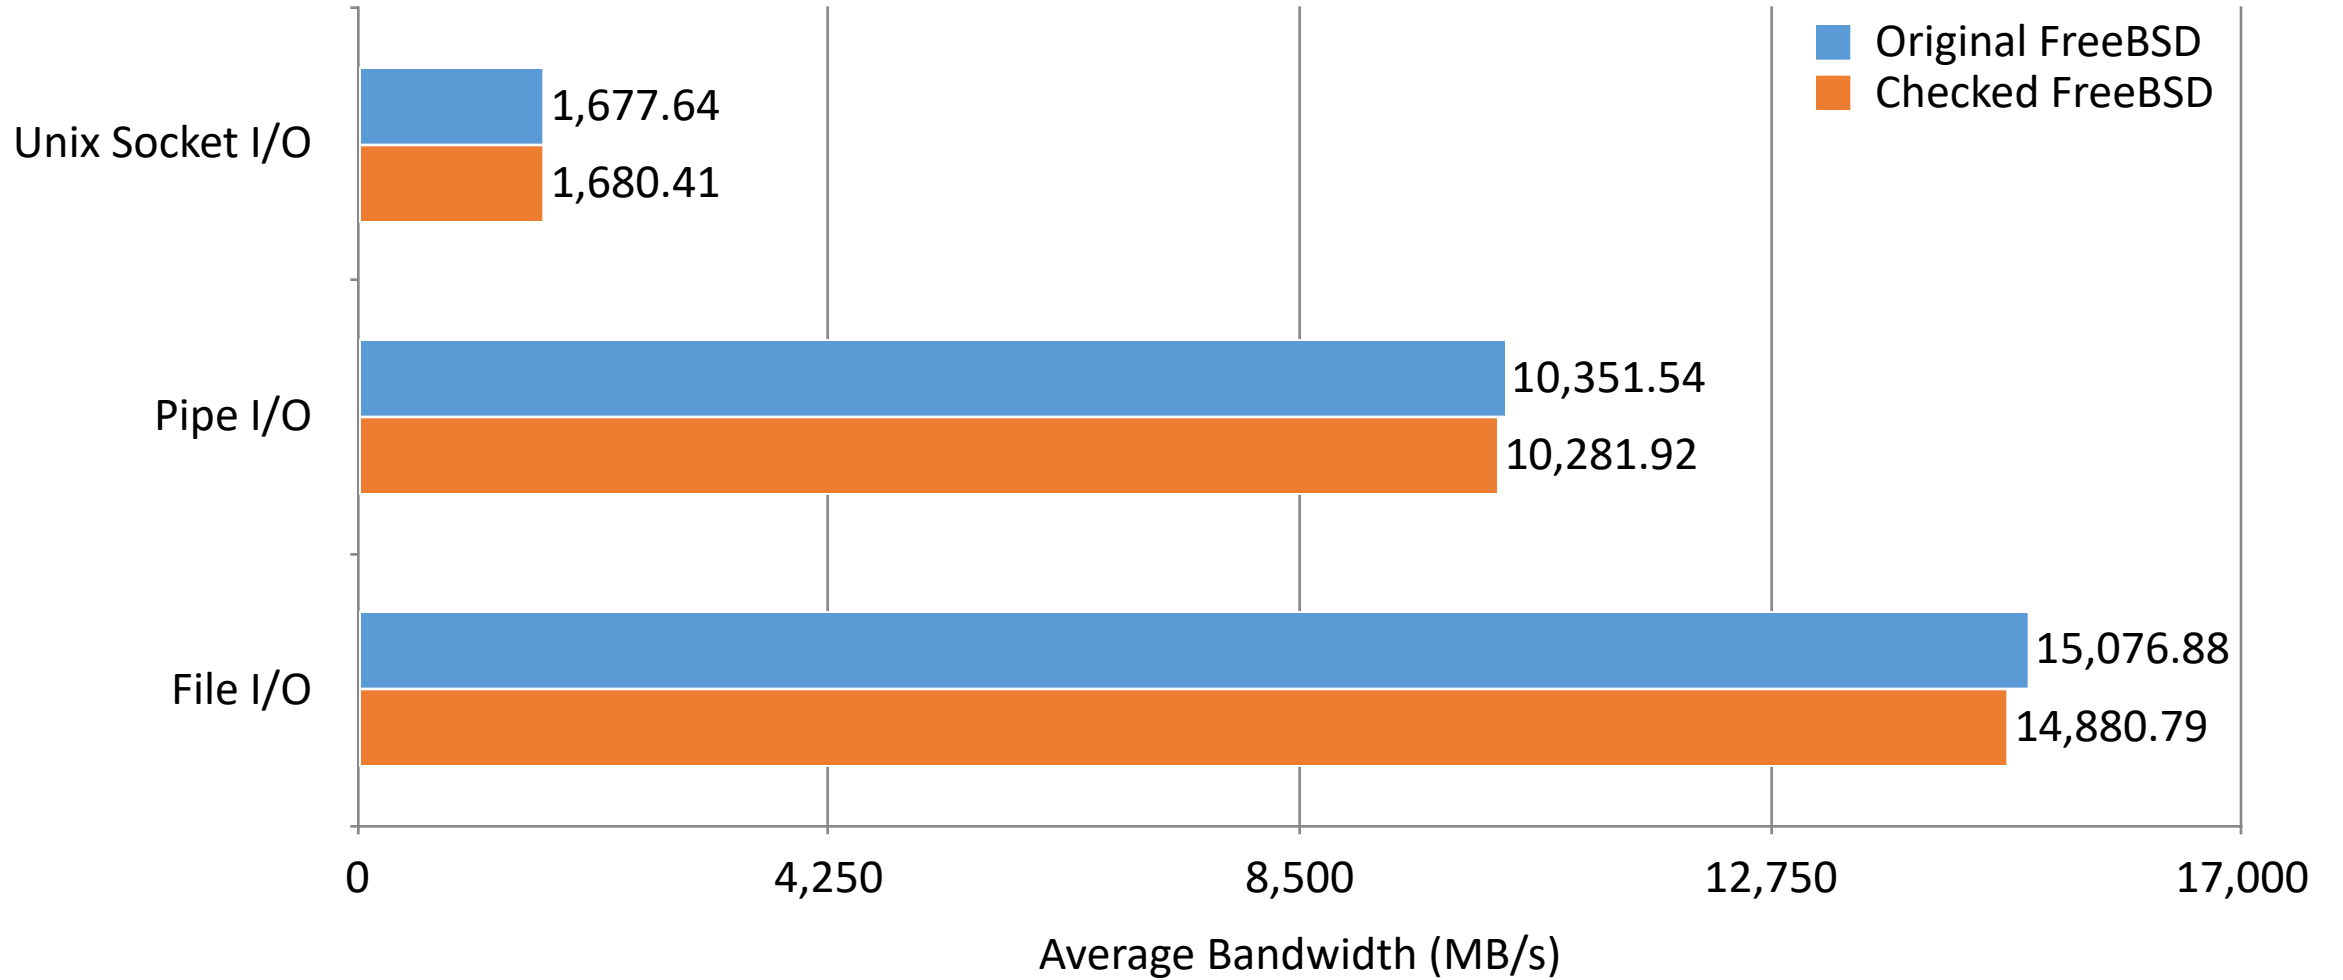
Evaluation - Part I

- System calls performance: tested using LMBench suite
- Test of subsystems
 - File System: the bandwidth for file read
 - System Calls: the latency of common system calls
 - Pipe: the bandwidth and transaction latency of pipe I/O
 - Unix Socket: the bandwidth and transaction latency for Unix domain sockets
- Test of system calls:
 - copyin() and copyout(): read(), write(), send(), recv()
 - copyinstr(): open()

Performance - System Call Latency Test



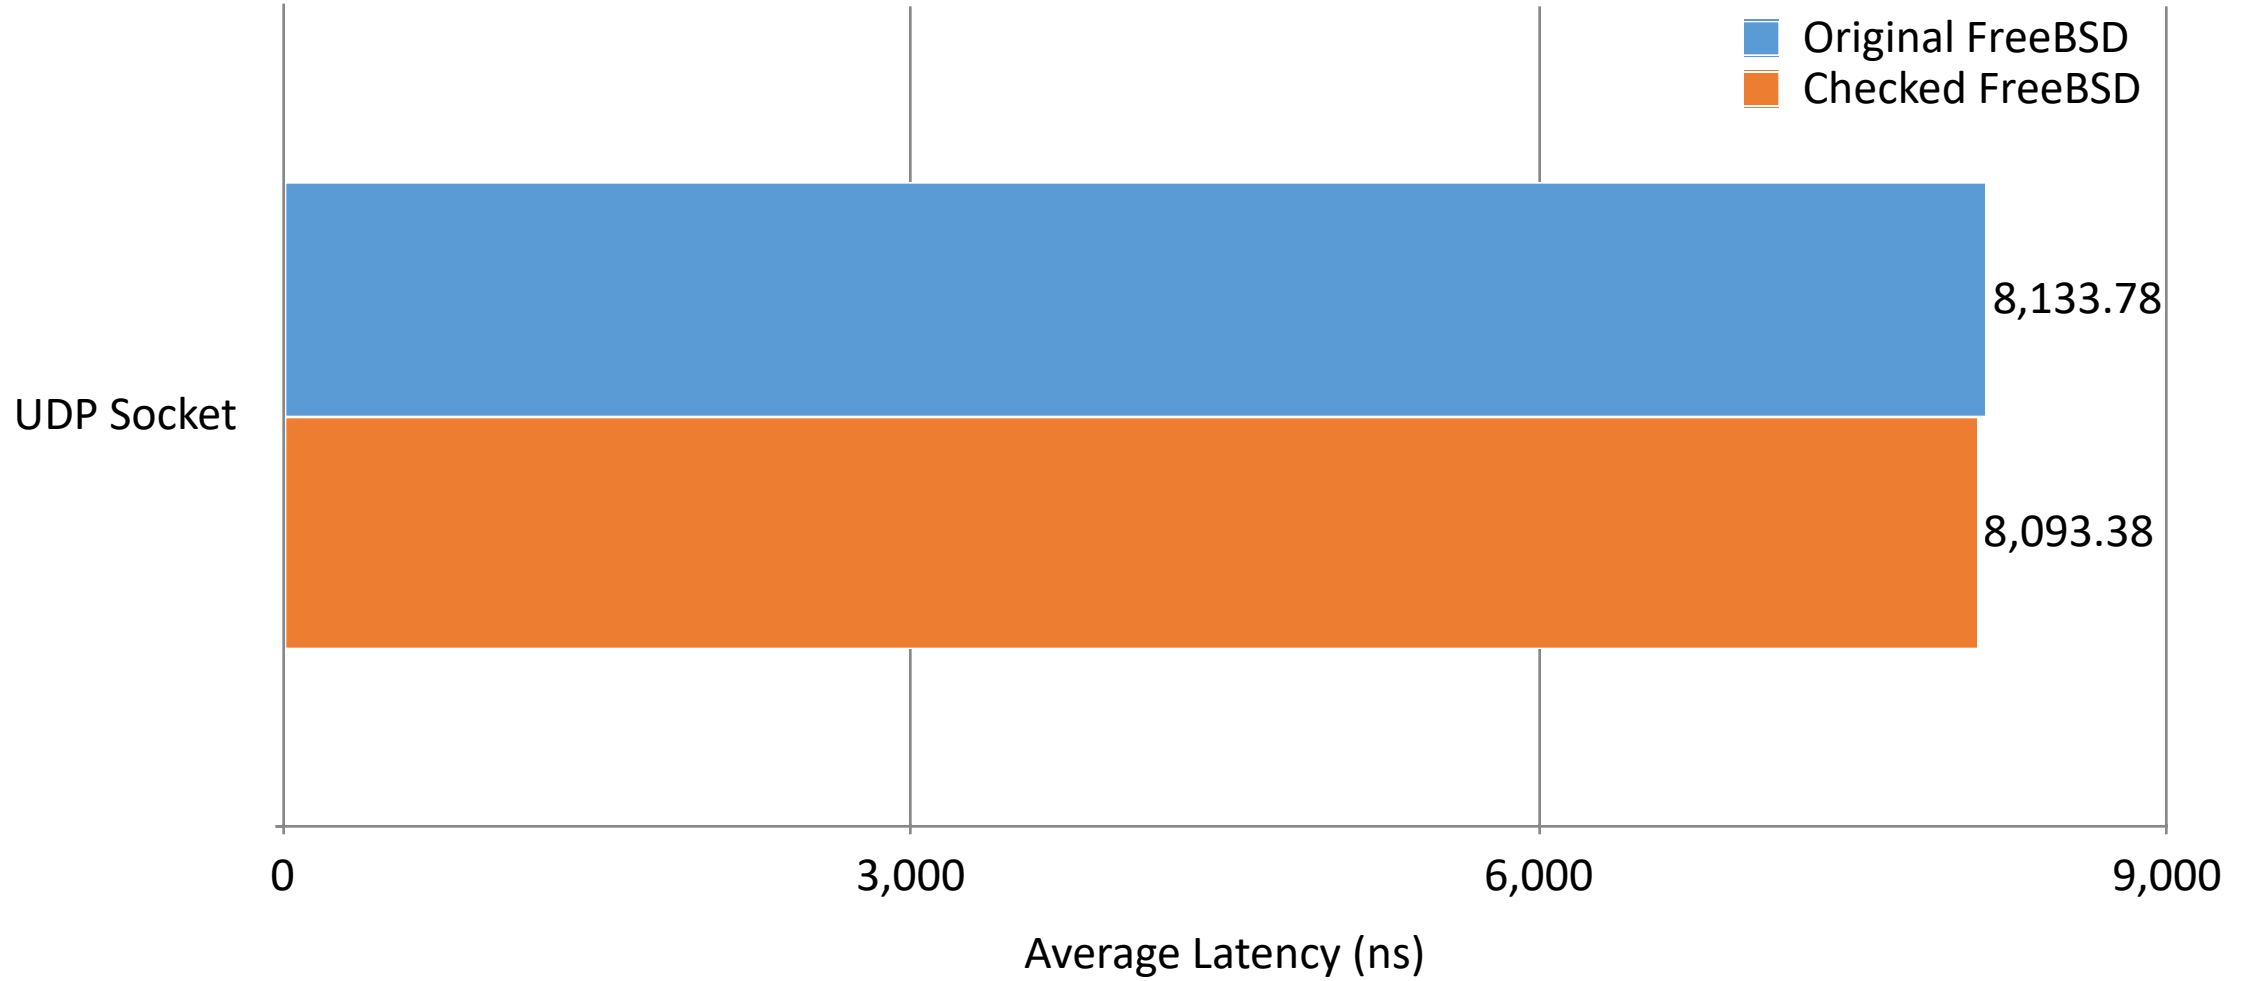
Performance - System Call Bandwidth Test



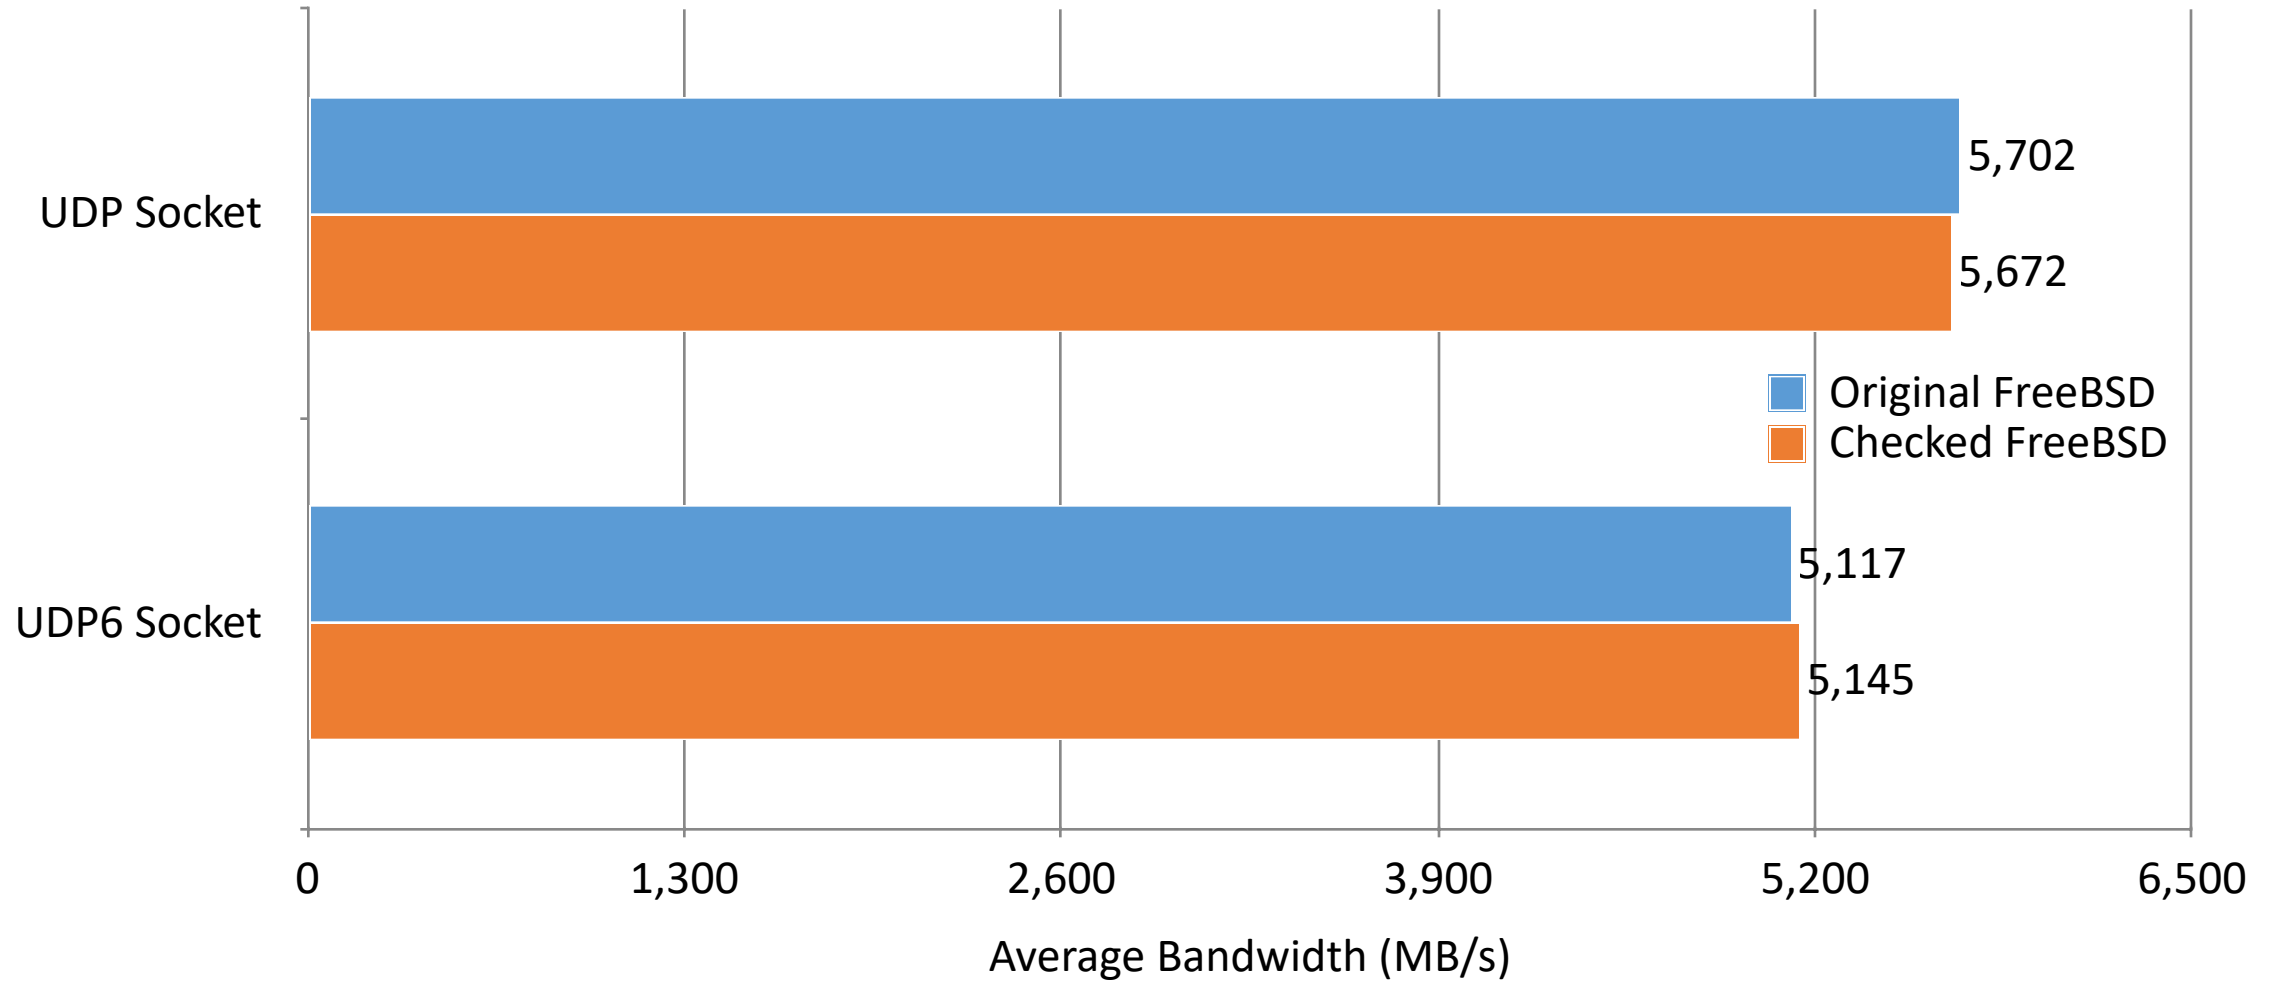
Evaluation - Part II

- UDP performance:
 - We used LMBench suite and iPerf3 to test UDP performance
 - LMBench: measured UDP latency
 - iPerf3: measured bandwidth for its server process

Performance - UDP Latency Test



Performance - UDP Bandwidth Test



Evaluation - Code Size Overhead

FreeBSD (baseline)	Checked FreeBSD	Checked - Origin
31,295,952	31,295,448	504

Size of kernel in bytes

Conclusion

- Refactored FreeBSD kernel to use checked pointers
 - copyin() and copyout() calls made by system calls
 - Packet manipulation code in IP and UDP
- Spatial memory safety guaranteed for packet manipulation in IP
 - Thwart buffer overflows and buffer overreads
- **Practically no performance or memory overhead**
- Low refactoring cost
 - Usually does not need to refactor the whole program structure/logic
 - Less than 100 hours spent by two undergraduate students to complete this project