



# Clobber-NVM: Log Less, Re-execute More

*Yi Xu*, Joseph Izraelevitz, Steven Swanson



UC San Diego & University of Colorado, Boulder

Published on ASPLOS 2021

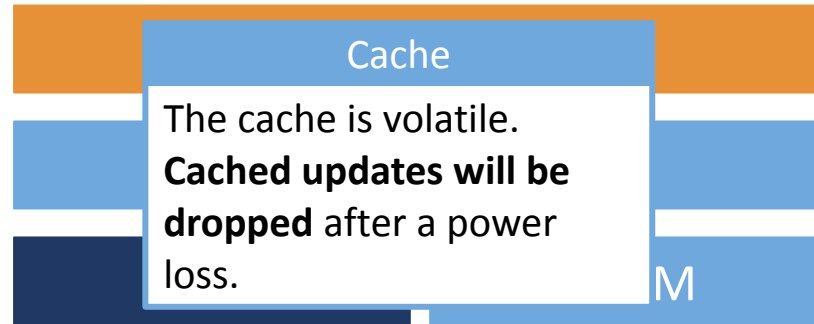
# Persistent Memory Programming



# Persistent Memory Programming



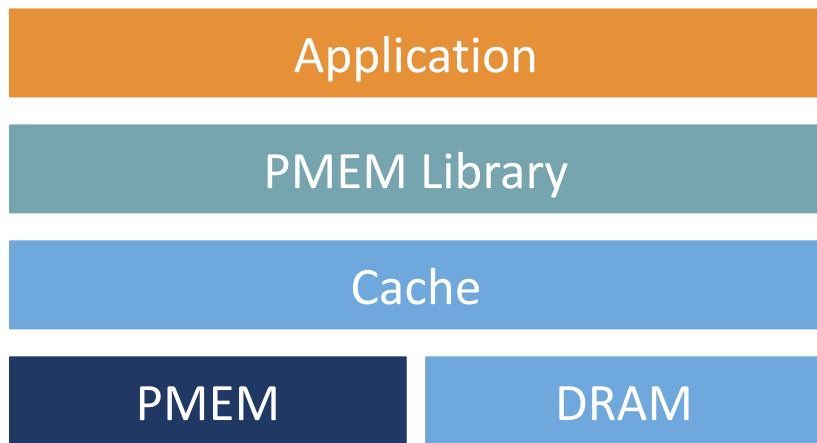
# Persistent Memory Programming



# Persistent Memory Programming



# Persistent Memory Programming



## PMEM library

- PMEM libraries provide the means to **apply sets of writes to persistent memory atomically.**
- Unfortunately, most current libraries impose **significant overhead.**

# PMEM Program with Undo Logging

## DRAM program

```
void list_push(list_t *list, char* value){
    memcpy(list->buf[list->size], value, strlen(value));
    list->size++;
}
```

## PMEM program with undo logging

```
void list_push(list_t *list, char* value){
    undo_log(value, strlen(value));
    persist_barrier();
    memcpy(list->buf[list->size], value, strlen(value));
    undo_log(list->size, sizeof(size_t));
    persist_barrier();
    list->size++;
}
```

# PMEM Program with Undo Logging

## DRAM program

```
void list_push(list_t *list, char* value){
    memcpy(list->buf[list->size], value, strlen(value));
    list->size++;
}
```

## PMEM program with undo logging

```
void list_push(list_t *list, char* value){
    undo_log(value, strlen(value));
    persist_barrier();
    memcpy(list->buf[list->size], value, strlen(value));
    undo_log(list->size, sizeof(size_t));
    persist_barrier();
    list->size++;
}
```



# PMEM Program with Undo Logging

## DRAM program

```
void list_push(list_t *list, char* value){
    memcpy(list->buf[list->size], value, strlen(value));
    list->size++;
}
```

## PMEM program with undo logging

```
void list_push(list_t *list, char* value){
    undo_log(value, strlen(value));
    persist_barrier();
    memcpy(list->buf[list->size], value, strlen(value));
    undo_log(list->size, sizeof(size_t));
    persist_barrier();
    list->size++;
}
```

# PMEM Program with Undo Logging

## DRAM program

```
void list_push(list_t *list, char* value){  
    memcpy(list->buf[list->size], value, strlen(value));  
    list->size++;  
}
```

## PMEM program with undo logging


```
void list_push(list_t *list, char* value){  
    undo_log(list, value);  
    persist(list, value);  
    memcpy(list->buf[list->size], value, strlen(value));  
    undo_log(list, value);  
    persist_barrier();  
    list->size++;  
}
```

Are the logs and barriers  
necessary?


# Recover Through Re-execution




# Recover Through Re-execution




# Recover Through Re-execution



# Recover Through Re-execution



# Recover Through Re-execution




# Recover Through Re-execution







# Recover Through Re-execution




# Recover Through Re-execution



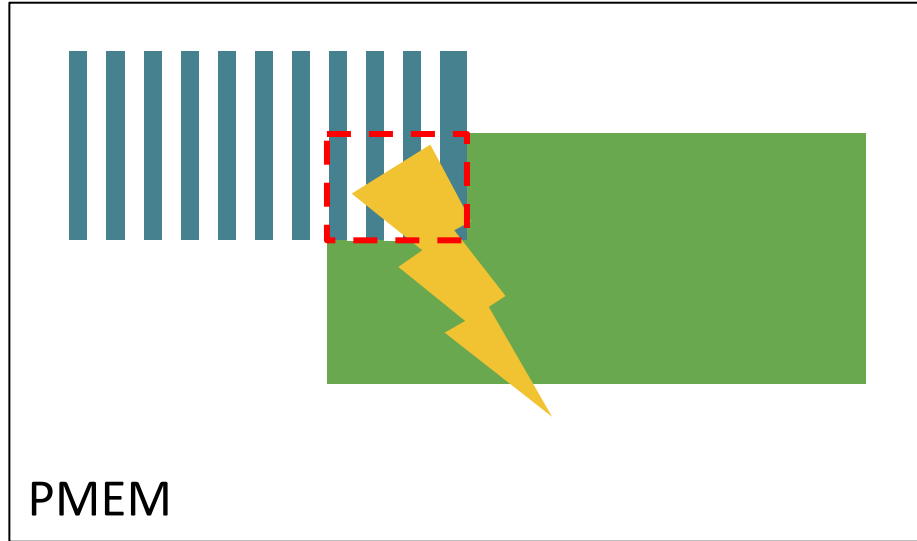
# Recover Through Re-execution



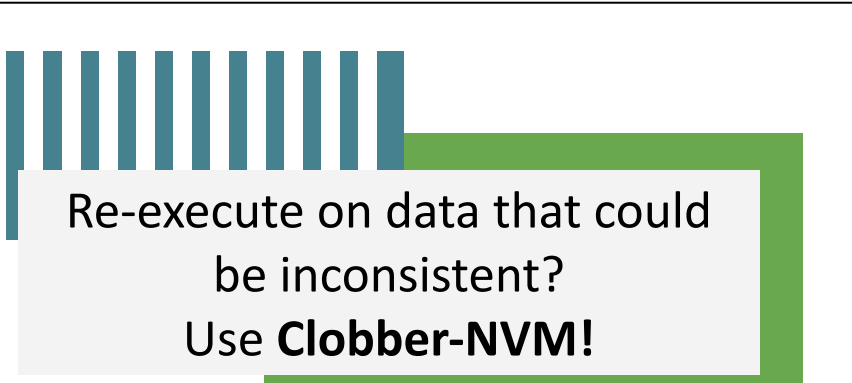
# Recover Through Re-execution



# Recover Through Re-execution



# Recover Through Re-execution



Re-execute on data that could  
be inconsistent?  
Use **Clobber-NVM!**

PMEM

The diagram shows a horizontal bar representing memory. The left portion consists of 12 vertical teal bars of equal height. The right portion is a solid green block that is taller than the teal bars. A light gray rectangular box is overlaid on the teal bars, containing the text 'Re-execute on data that could be inconsistent? Use Clobber-NVM!'. The word 'PMEM' is written in black text at the bottom left of the diagram's border.

# Clobber-NVM Program

## PMEM program with Clobber-NVM

```
void list_push(list_t *list, char* value){  
    txbegin();  
    memcpy(list->buf[list->size], value, strlen(value);  
    list->size++;  
    txend();  
}
```

Transaction  
Boundary

# Clobber-NVM Program

## PMEM program with Clobber-NVM

```
void list_push(list_t *list, char* value){  
    txbegin();  
    memcpy(list->buf[list->size], value, strlen(value);  
    list->size++;  
    txend();  
}
```

Transaction  
Boundary

Input



# Clobber-NVM Program

## PMEM program with Clobber-NVM

```
void list_push(list_t *list, char* value){  
    txbegin();  
    memcpy(list->buf[list->size], value, strlen(value);  
    list->size++;  
    txend();  
}
```

Transaction  
Boundary

Input

Output

# Clobber-NVM Program

## PMEM program with Clobber-NVM

```
void list_push(list_t *list, char* value){  
    txbegin();  
    memcpy(list->buf[list->size], value, strlen(value);  
    list->size++;  
    txend();  
}
```

Transaction  
Boundary

Input

Output

Input &  
Output

# Clobber-NVM Program

## PMEM program with Clobber-NVM

```
void list_push(list_t *list, char* value){  
    txbegin();  
    memcpy(list->buf[list->size], value, strlen(value);  
    list->size++;  
    txend();  
}
```

Transaction  
Boundary

Input

Output

Clobber  
Input

# Clobber-NVM Program

## PMEM program with Clobber-NVM

```
void list_push(list_t *list, char* value){  
    txbegin();  
    memcpy(list->buf[list->size], value, strlen(value);  
    list->size++;  
    txend();  
}
```

A transaction input is a clobbered input if it may be overwritten within this transaction, and this write is a clobber write.

Transaction  
Boundary


Input

Output

Clobber  
Input


Clobber  
Write

# Clobbered Inputs




**Clobbered inputs are a problem for re-execution.**

# Clobbered Inputs




**Clobbered inputs are a problem for re-execution.**

# Clobbered Inputs



**Clobbered inputs are a problem for re-execution.**


# Clobber\_Log before Clobber Writes




**Clobber\_Log --- undo logs before clobber writes.**




# Clobber\_log before Clobber Writes



# Clobber\_log before Clobber Writes



# Re-execute Based on Clobber\_log



# Re-execute Based on Clobber\_log




# Re-execute Based on Clobber\_log




# Re-execute Based on Clobber\_log



# Re-execute Based on Clobber\_log




# Handle DRAM Accesses






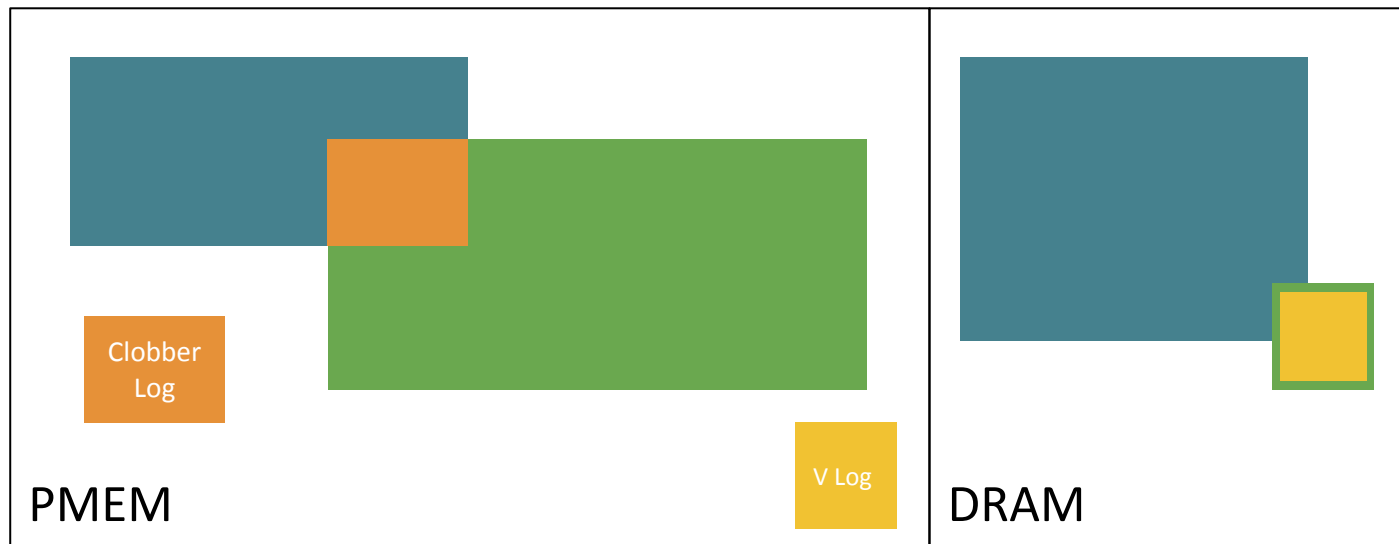
# Handle DRAM Accesses



# Handle DRAM Accesses




# Handle DRAM Accesses on v\_log




**v\_log stores volatile transaction inputs**


# Handle DRAM Accesses




# Handle DRAM Accesses



# Handle DRAM Accesses



# Handle DRAM Accesses

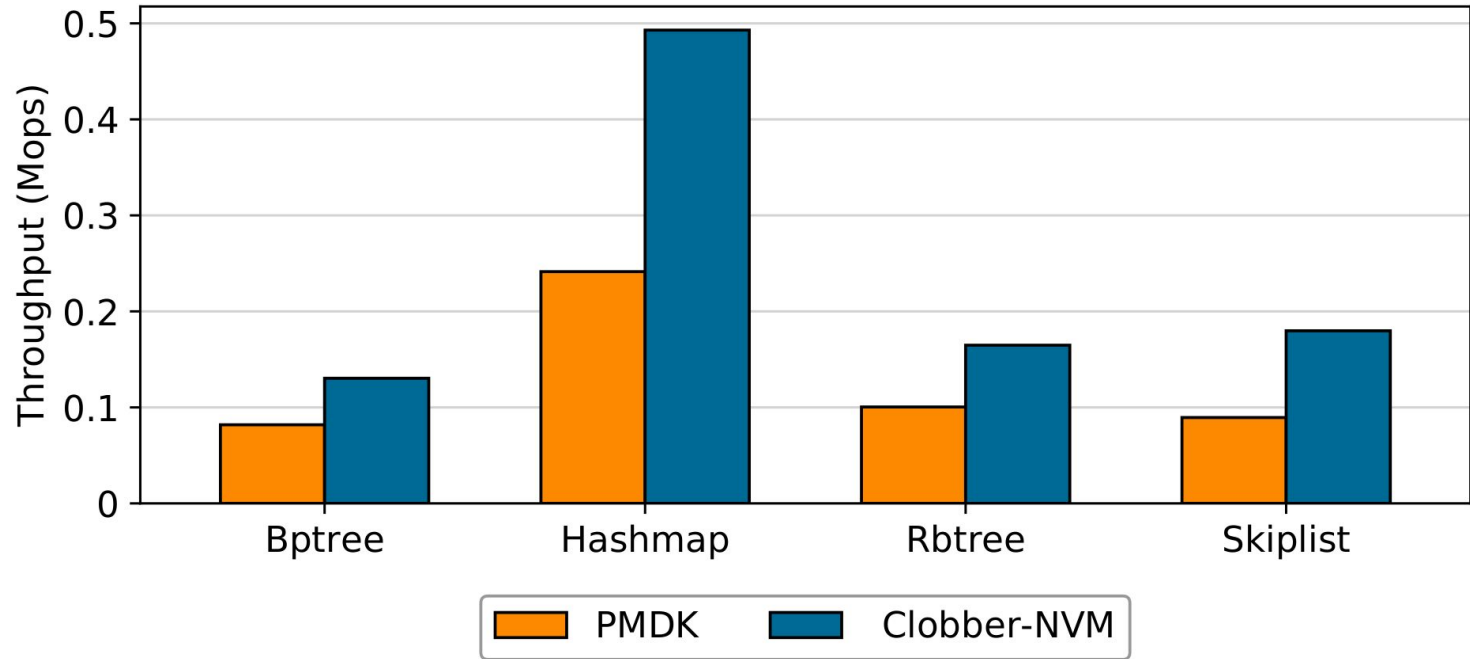


# Evaluation Setup


- Platform: two 24-core Intel Cascade Lake SP processors, running at 2.2 GHz. The platform has a total of 192 GB of DRAM and 1.5 TB (6 ×256 GB) of Intel Optane DC Persistent Memory directly attached to each processor.
- Configured Optane DCPMM in 100% App Direct mode.
- All experiments use Ext4 to manage persistent pools and directly access NVM pages via DAX.



# Data structure Benchmarks



# Memcached Performance



# Conclusion

- Clobber-NVM: Recovers by re-executing interrupted transactions.
- Clobber-NVM compiler: Identifies necessary log entries, and automatically adds logging for selected variables.
- Evaluation shows that Clobber-NVM has high performance.