




HOW WILL PERSISTENT MEMORY CHANGE SOFTWARE DESIGN

Maciej Maciejewski

HARDWARE



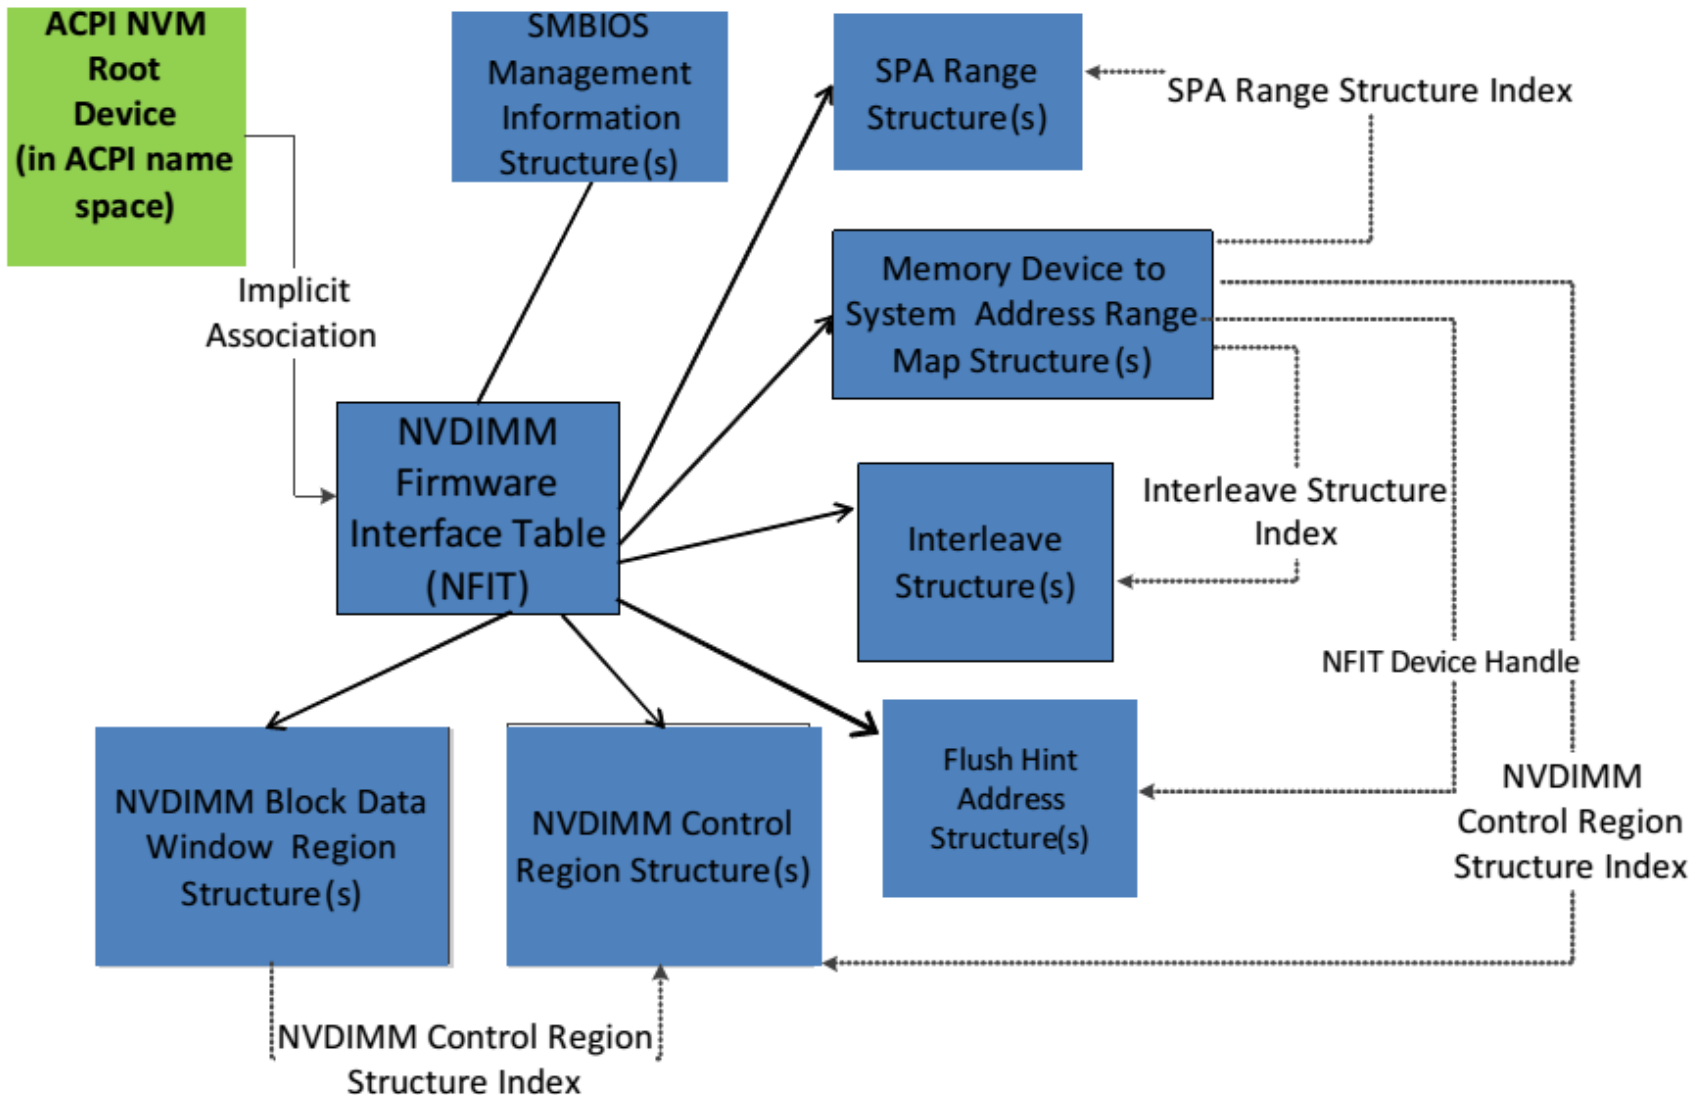
Persistent memory

How does it differ from DRAM

How fast is it?

OPERATING SYSTEM

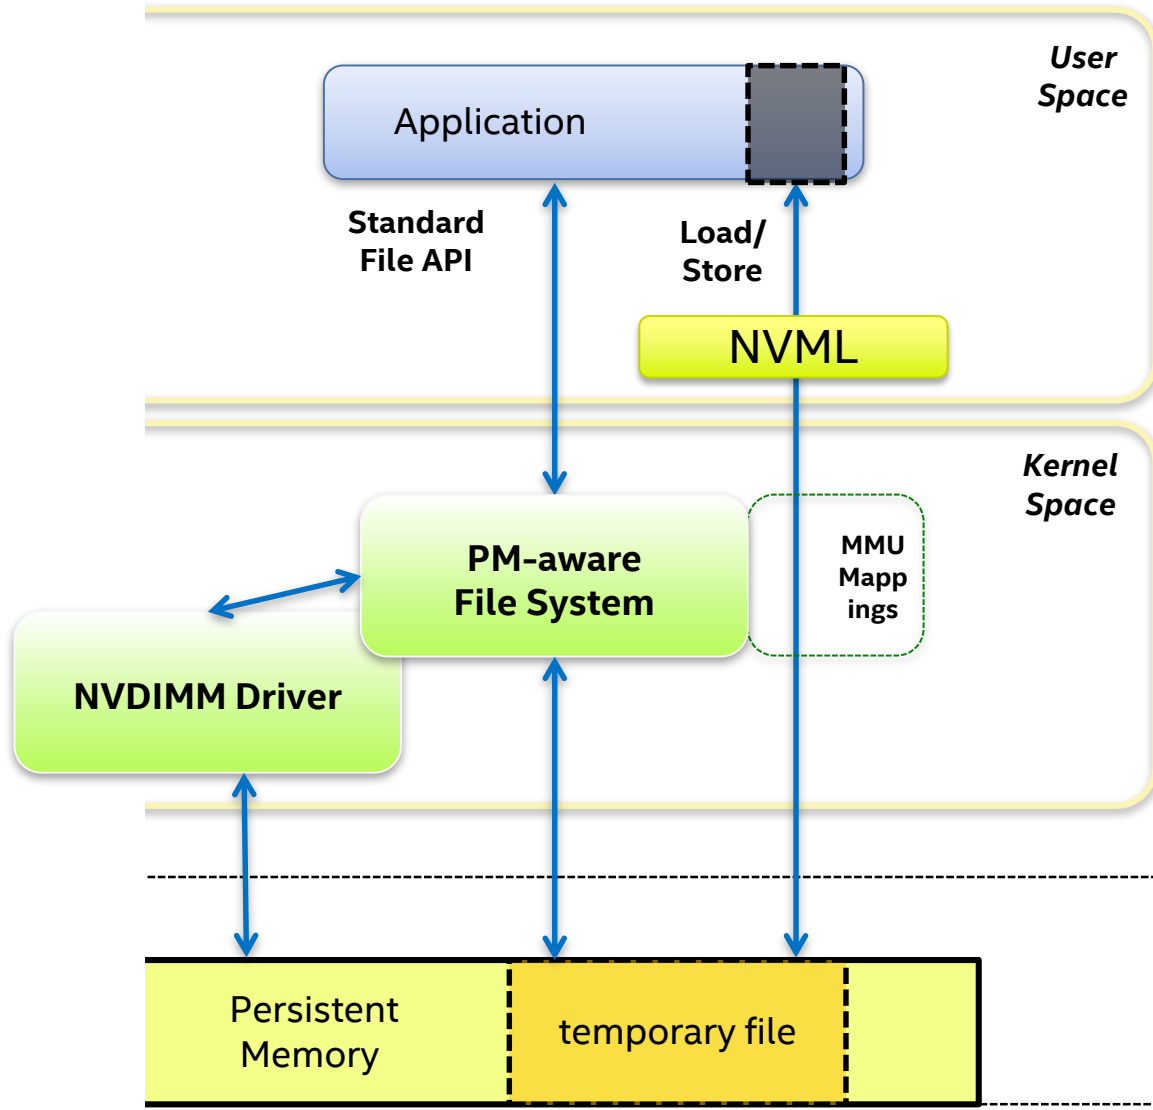
ACPI NFIT



E820

```
0.000000] NX (Execute Disable) protection: active
0.000000] e820: user-defined physical RAM map:
0.000000] user: [mem 0x0000000000000000-0x000000000009dbff] usable
0.000000] user: [mem 0x000000000009dc00-0x000000000009ffff] reserved
0.000000] user: [mem 0x00000000000e0000-0x00000000000fffff] reserved
0.000000] user: [mem 0x0000000000100000-0x0000000000d9ff9fff] usable
0.000000] user: [mem 0x00000000d9ffa000-0x00000000da00ffff] ACPI NVS
0.000000] user: [mem 0x00000000da001000-0x00000000da44dfff] usable
0.000000] user: [mem 0x00000000da44e000-0x00000000da7f9fff] reserved
0.000000] user: [mem 0x00000000da7fa000-0x00000000dad9efff] usable
0.000000] user: [mem 0x00000000dad9f000-0x00000000dae26fff] ACPI NVS
0.000000] user: [mem 0x00000000dae27000-0x00000000dae54fff] usable
0.000000] user: [mem 0x00000000dae55000-0x00000000daf67fff] reserved
0.000000] user: [mem 0x00000000daf68000-0x00000000dafc3fff] usable
0.000000] user: [mem 0x00000000dafc4000-0x00000000db044fff] reserved
0.000000] user: [mem 0x00000000db045000-0x00000000db04afff] usable
0.000000] user: [mem 0x00000000db04b000-0x00000000db062fff] ACPI data
0.000000] user: [mem 0x00000000db063000-0x00000000db0b9fff] usable
0.000000] user: [mem 0x00000000db0ba000-0x00000000db0c6fff] reserved
0.000000] user: [mem 0x00000000db0c7000-0x00000000dbe64fff] usable
0.000000] user: [mem 0x00000000dbe65000-0x00000000dbea4fff] reserved
0.000000] user: [mem 0x00000000dbea5000-0x00000000dbef3fff] usable
0.000000] user: [mem 0x00000000dbef4000-0x00000000dbf23fff] reserved
0.000000] user: [mem 0x00000000dbf24000-0x00000000dbf26fff] ACPI data
0.000000] user: [mem 0x00000000dbf27000-0x00000000dbf2efff] ACPI NVS
0.000000] user: [mem 0x00000000dbf2f000-0x00000000dbf3dfff] reserved
0.000000] user: [mem 0x00000000dbf3e000-0x00000000dbf4efff] ACPI NVS
0.000000] user: [mem 0x00000000dbf4f000-0x00000000dbf4ffff] reserved
0.000000] user: [mem 0x00000000dbf50000-0x00000000dbf60fff] ACPI NVS
0.000000] user: [mem 0x00000000dbf61000-0x00000000dbf8cfff] reserved
0.000000] user: [mem 0x00000000dbf8d000-0x00000000dbf93fff] ACPI NVS
0.000000] user: [mem 0x00000000dbf94000-0x00000000dbfe3fff] reserved
0.000000] user: [mem 0x00000000dbfe4000-0x00000000dbffcfff] usable
0.000000] user: [mem 0x00000000dbffd000-0x00000000dbffefff] reserved
0.000000] user: [mem 0x00000000dbfff000-0x00000000dbffffff] usable
0.000000] user: [mem 0x00000000dd000000-0x00000000df1fffff] reserved
0.000000] user: [mem 0x00000000df800000-0x00000000fbffffff] reserved
0.000000] user: [mem 0x00000000fec00000-0x00000000fec00fff] reserved
0.000000] user: [mem 0x00000000fed00000-0x00000000fed03fff] reserved
0.000000] user: [mem 0x00000000fed1c000-0x00000000fed1ffff] reserved
0.000000] user: [mem 0x00000000fee00000-0x00000000fee00fff] reserved
0.000000] user: [mem 0x00000000ff000000-0x00000000ffffffff] reserved
0.000000] user: [mem 0x0000000100000000-0x000000027fffffff] usable
0.000000] user: [mem 0x0000000300000000-0x000000033fffffff] persistent (type 12)
0.000000] user: [mem 0x0000000340000000-0x000000037fffffff] persistent (type 12)
0.000000] user: [mem 0x0000000380000000-0x00000003bfffffff] persistent (type 12)
0.000000] user: [mem 0x00000003c0000000-0x000000041cdfffff] usable
0.000000] SMBIOS 2.7 present
```


WHAT CAN WE DO WITH IT?



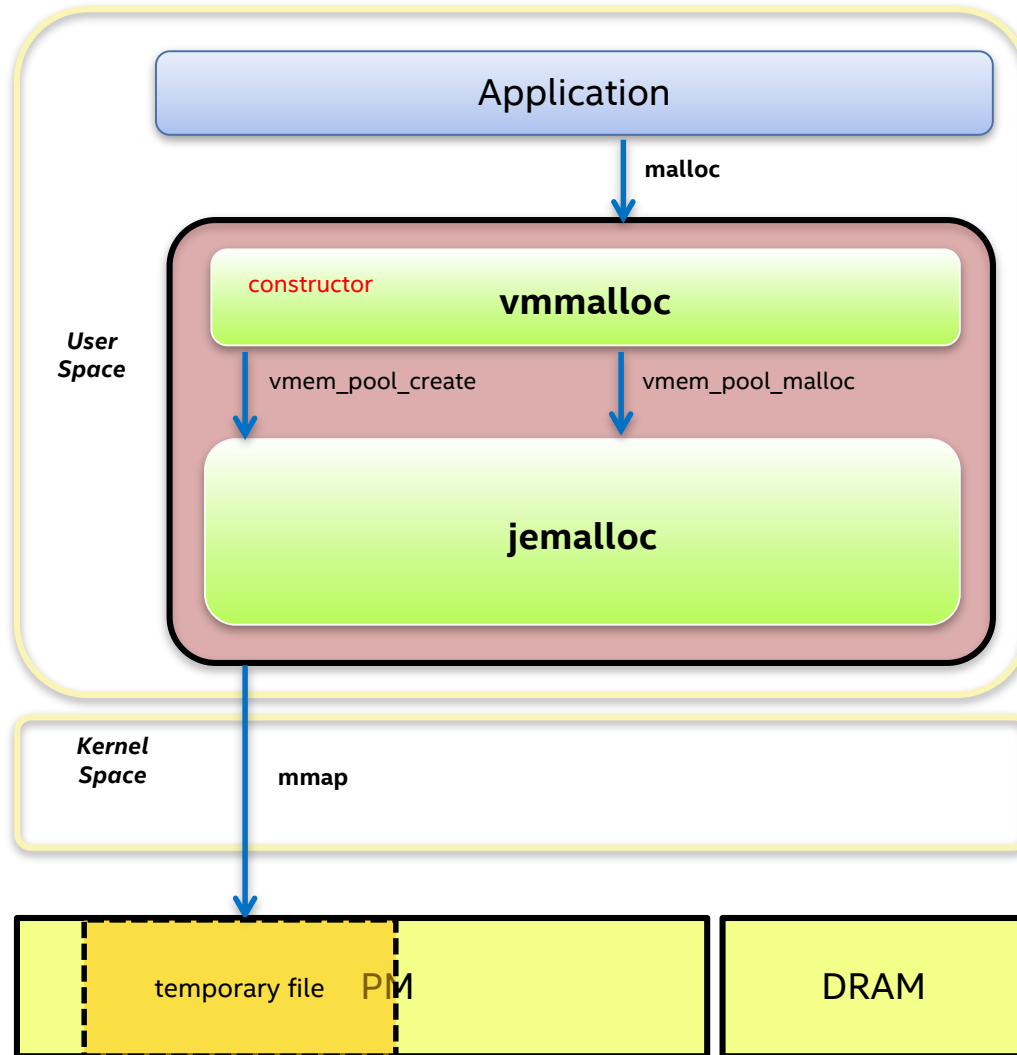
NVML

<http://pmem.io>

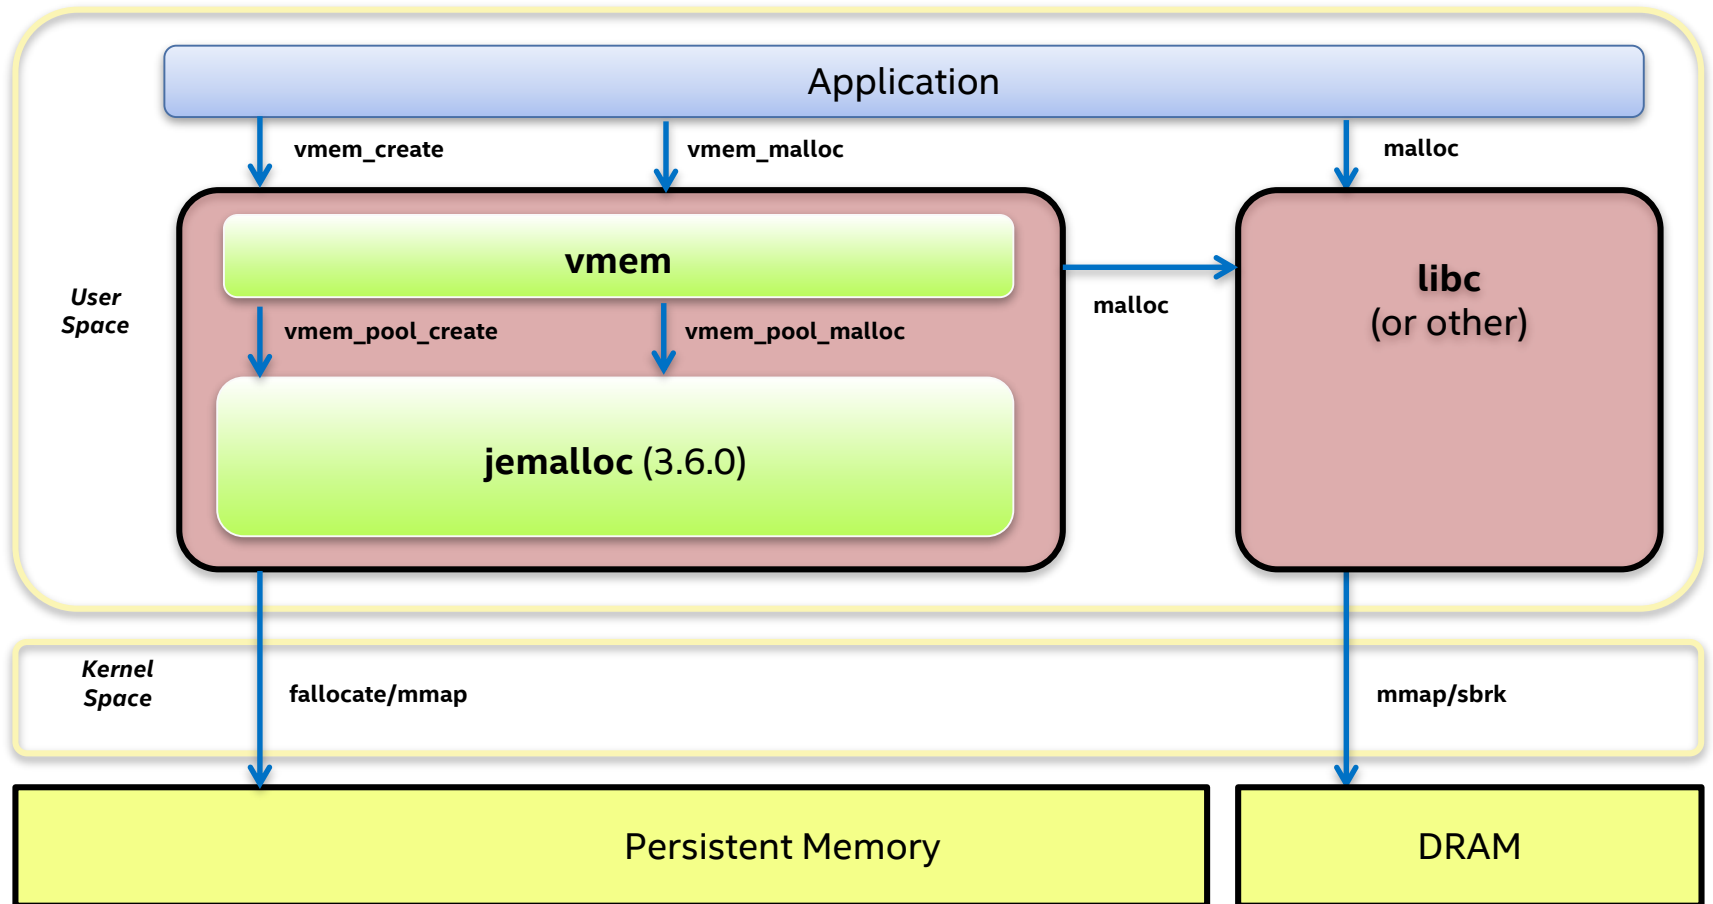
<https://github.com/pmem/nvml/>

Memory

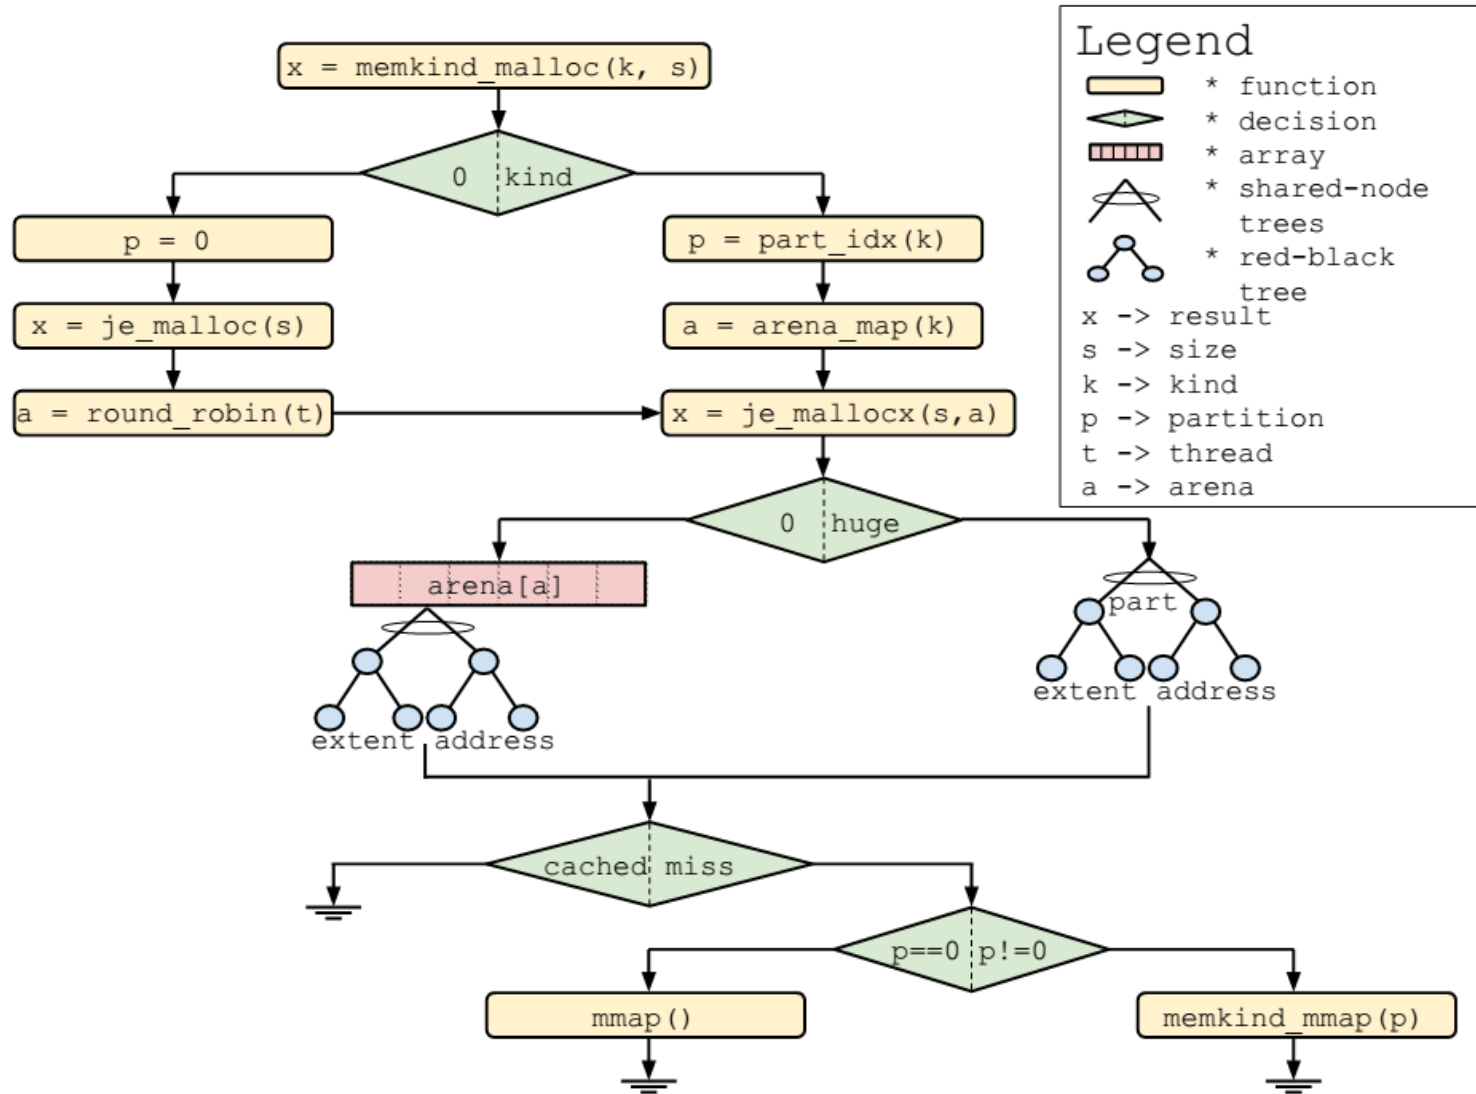
Memory - libvmmalloc




Memory - libvmem



Memory - memkind

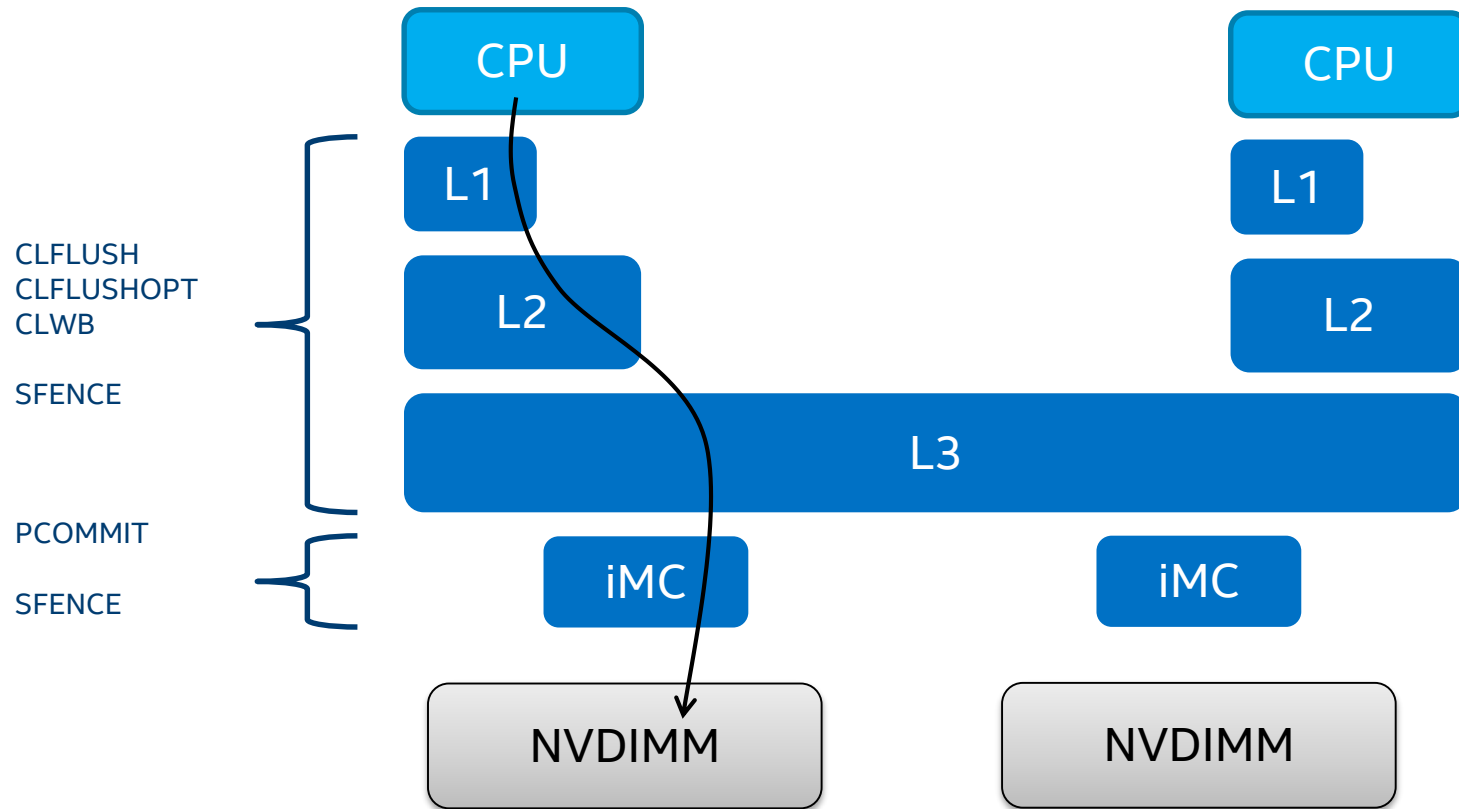


Block storage




Byte persistency

Why persistency is an issue?



Flushing to Persistence

```
open(...);  
mmap(...);  
strcpy(pmem, „Hello”);  
msync(pmem, 6, MS_SYNC);  
pmem_persist(pmem, 6);  
strcpy(pmem, „Hello, World!”);  
pmem_persist(pmem, 14);
```



Crossing the 8-Byte Store

Result?

1. „\0\0\0\0\0\0\0\0\0\0...”
2. „Hello, w\0\0\0\0\0\0...”
3. „\0\0\0\0\0\0\0\0\0orld!\0”
4. „Hello, \0\0\0\0\0\0\0\0”
5. „Hello, World!\0”

nvml Persistent Libraries

libpmem – Basic persistency handling

libpmemblk – Block access to pmem

libpmemlog - Log file on pmem (append-mostly)

libpmemobj - Transactional Object Store on pmem

libpmemobj

Object API

Transactional API


```
20
21  /* create a new collection */
22  TX_BEGIN(pop) {
23      tree_map_new(pop, &D_RW(root)->map);
24
25      /*
26       * We don't want an empty collection, so let's insert
27       * a few new entries in the same transaction.
28       */
29      TOID(struct store_item) n;
30
31      n = TX_NEW(struct store_item);
32      tree_map_insert(pop, D_RO(root)->map, 5, n.oid);
33
34      n = TX_NEW(struct store_item);
35      tree_map_insert(pop, D_RO(root)->map, 10, n.oid);
36  } TX_END
37
```

APPLICATIONS


pm_invaders

Redis key/value store


Startup time




DRAM usage




Running out of DRAM




Running out of DRAM




Running out of DRAM




Running out of DRAM




CPU usage



CPU usage

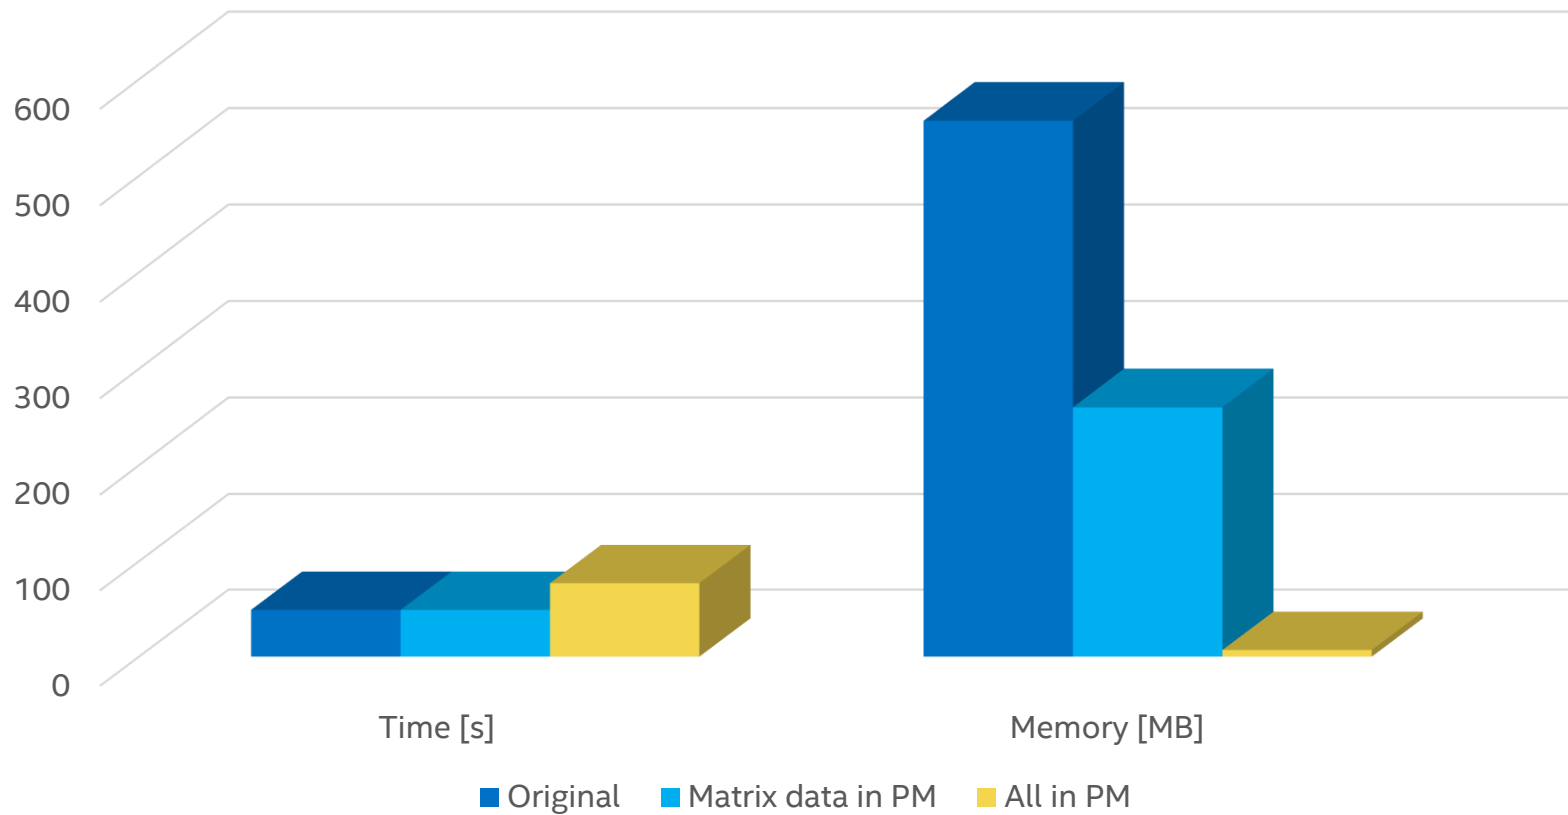


CPU usage



Moving even more data

GMRES Sparse Matrix solver



Modifying app engine

```
for (int i=0; i<NUMBER_OF_ITERATIONS; i++) {  
    result = calculateThis(i, result);  
}
```

```
i_pm = 0; //at first runtime of app only
```

```
...
```

```
...
```

```
for (i_pm; i_pm<NUMBER_OF_ITERATIONS; i_pm++) {  
    TX_BEGIN(pool) {  
        result_pm = calculateThis(i_pm, result_pm);  
    } TX_END  
}  
i_pm = 0;
```

Possibilities

Choices

Even more fun ...

