



## The Open Source DRAM Simulator DRAMSys

Prof. Dr. Matthias Jung, JMU Würzburg



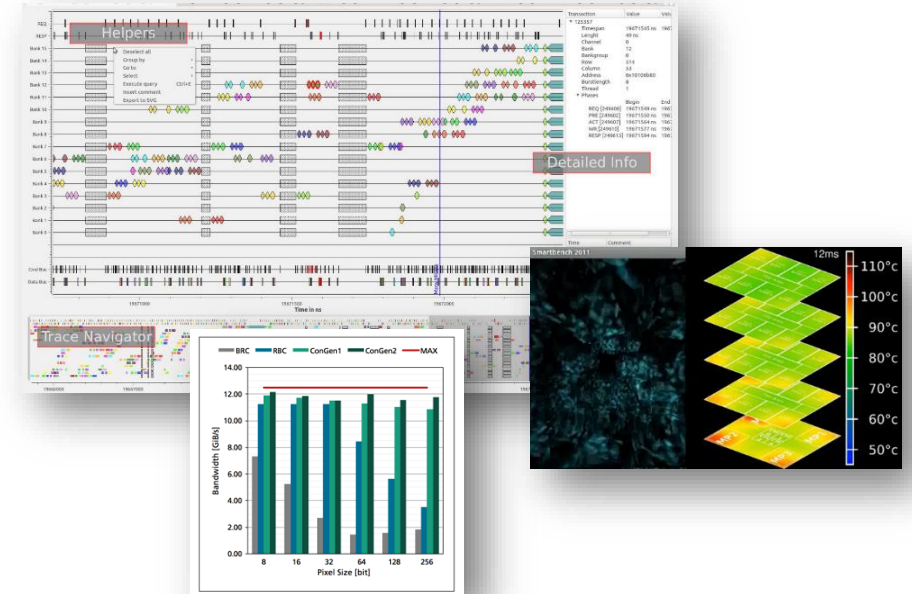
# DRAMSys in a Nutshell

## Simulation and Design Space Exploration of Modern DRAM-based Memory Systems:

- Which DRAM configuration?
- When to support DDR5 or LPDDR5?
- How to configure the memory controller?
- What is the system-level application behavior?

## DRAMSys Offers:

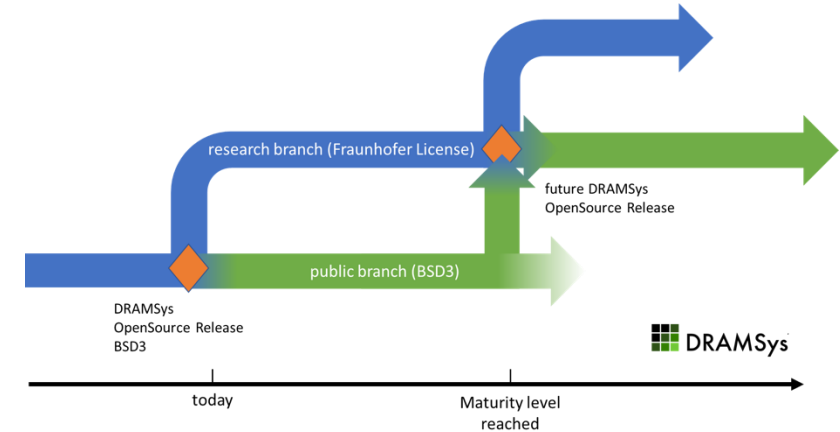
- High-speed and flexible models of all standards
- Fast and accurate design space exploration
- Early identification of bottlenecks
- Connection to cores (e.g. SystemC, gem5, ...)



# DRAMSys Open Source Model

Fork me on Github!

- Open source: DDR3/4, LPDDR4, Wide I/O 1/2, GDDR5/X, GDDR6, and HBM2
- Commercial/academic licenses: DDR5, LPDDR5, HBM3, Trace Analyzer tool
- New standard models will be open-sourced when a level of maturity is reached
- Customer-specific consulting, modifications and developments



Thanks to our Key Partners:

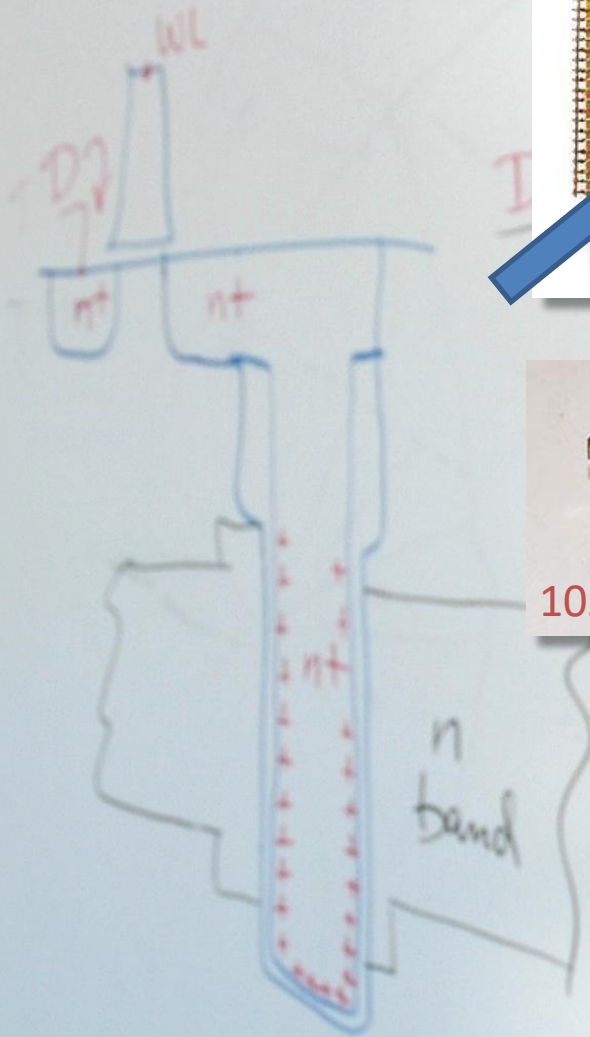


**Rambus**




Mercedes-Benz

**ARTERIS IP**





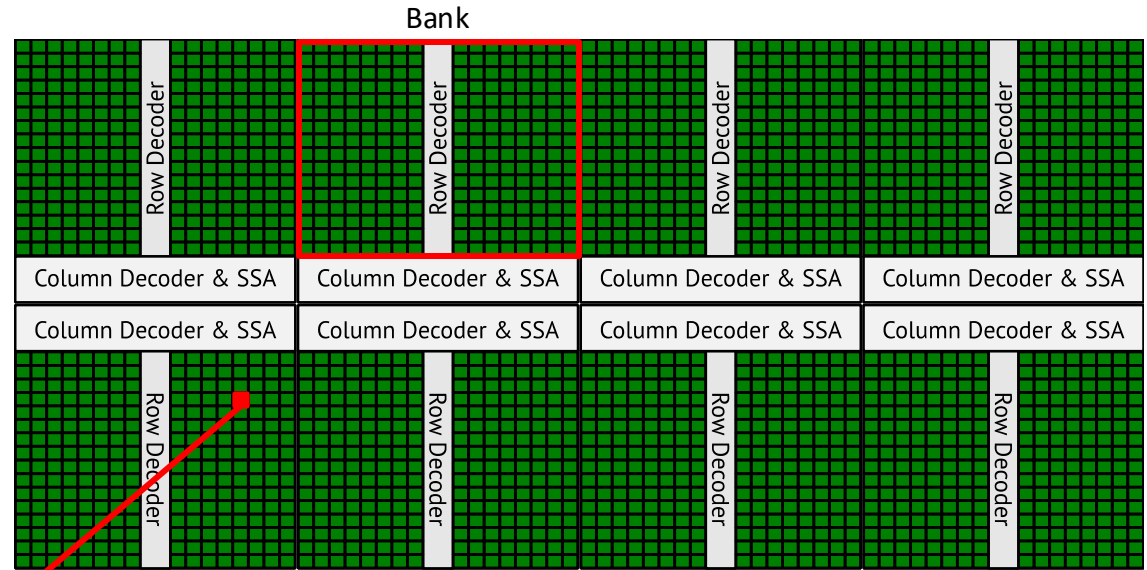


transistor  
capacitor



Invented by Robert H. Dennard in 1966




# The DRAM Device / Operation



- Using Sub-Arrays for efficient wiring
- Bank parallelism, but banks share data and command bus



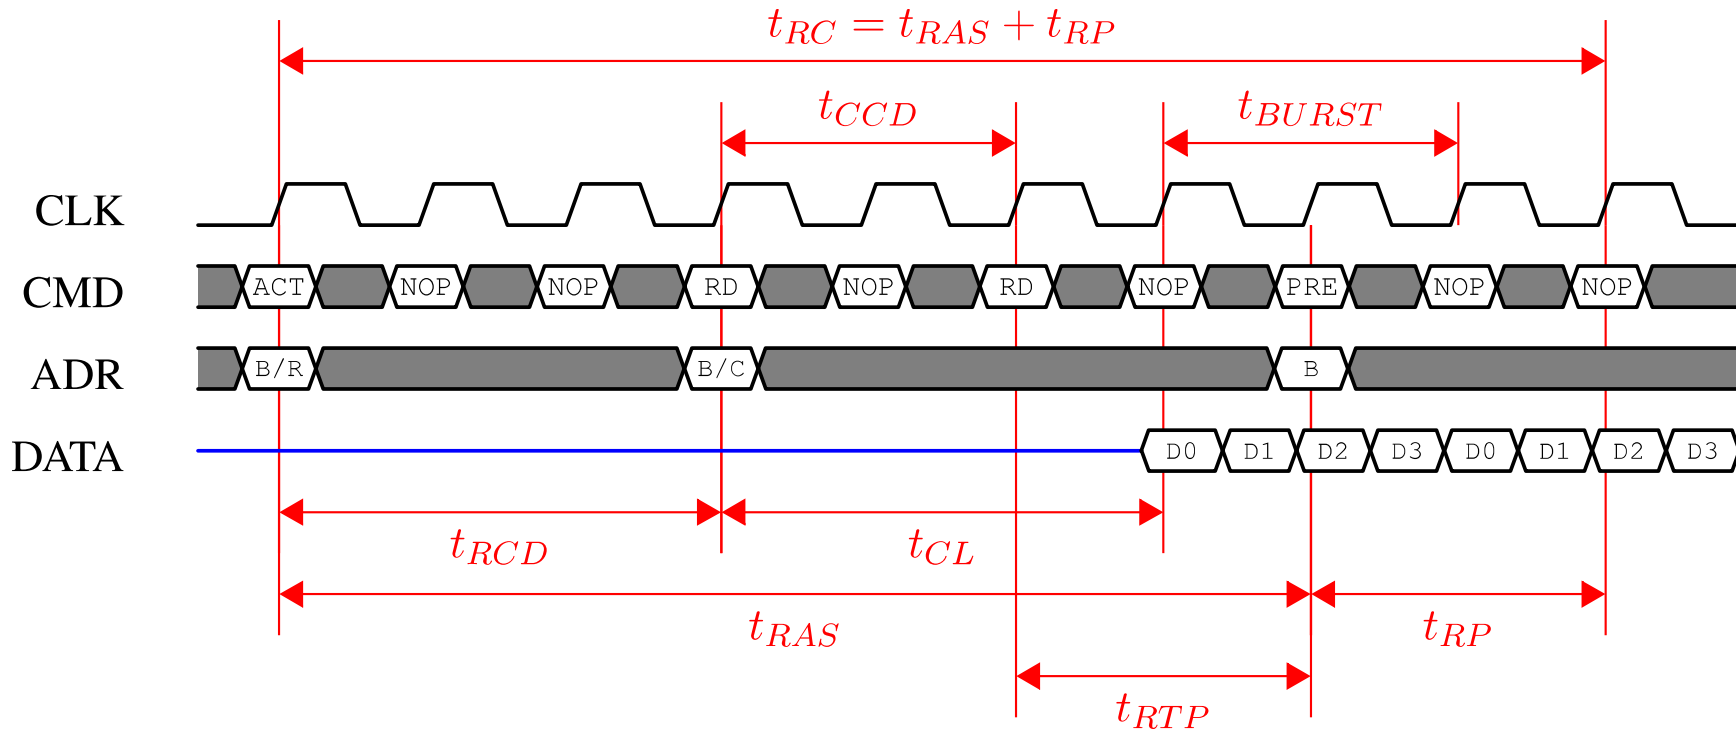
# DRAMs Basic Operations



## Important DRAM Commands:

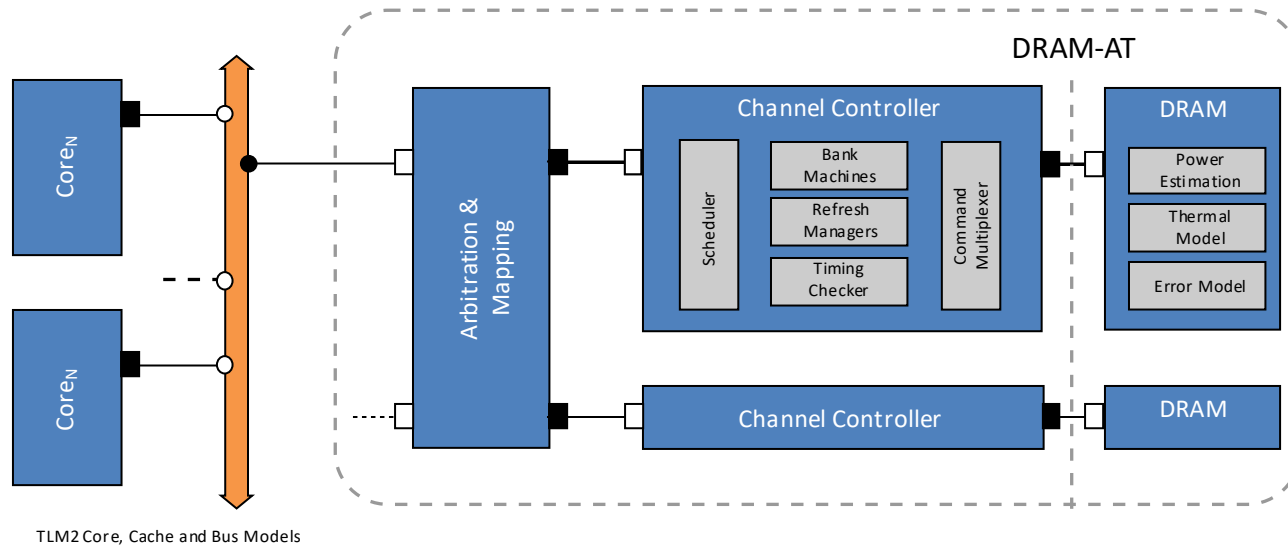
- **ACT:** Activates a specific row in a specific bank (sensing into PSA) [ $t_{RCD}$ ]
- **RD:** Read from activated row (prefetch from PSA to SSA and burst out) [ $t_{CL} + t_{BURST}$ ]
- **PRE:** Precharges set  $LWL=0$  set  $LBL=VDD/2$  [ $t_{RP}$ ]
- **REFA:** DRAM cells are leaky and have to be refreshed [ $t_{REFI}$  &  $t_{RFC}$ ]

# JEDEC Standard: e.g. Timing Dependencies



Timing dependencies must be fulfilled by the DRAM controller

# DRAMSys Architecture

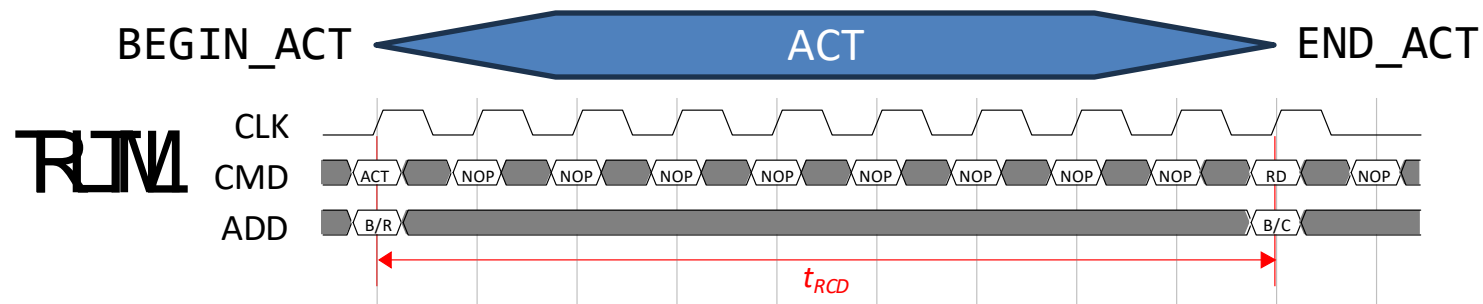


- Based on SystemC TLM2, compliant with TLM-AT coding style
- Flexible SW-Architecture to support various JEDEC DRAM standards (e.g., DDR4, LPDDR4, GDDR6, HBM, ...)
- For RTL-like accuracy a custom TLM protocol (DRAM-AT) is used



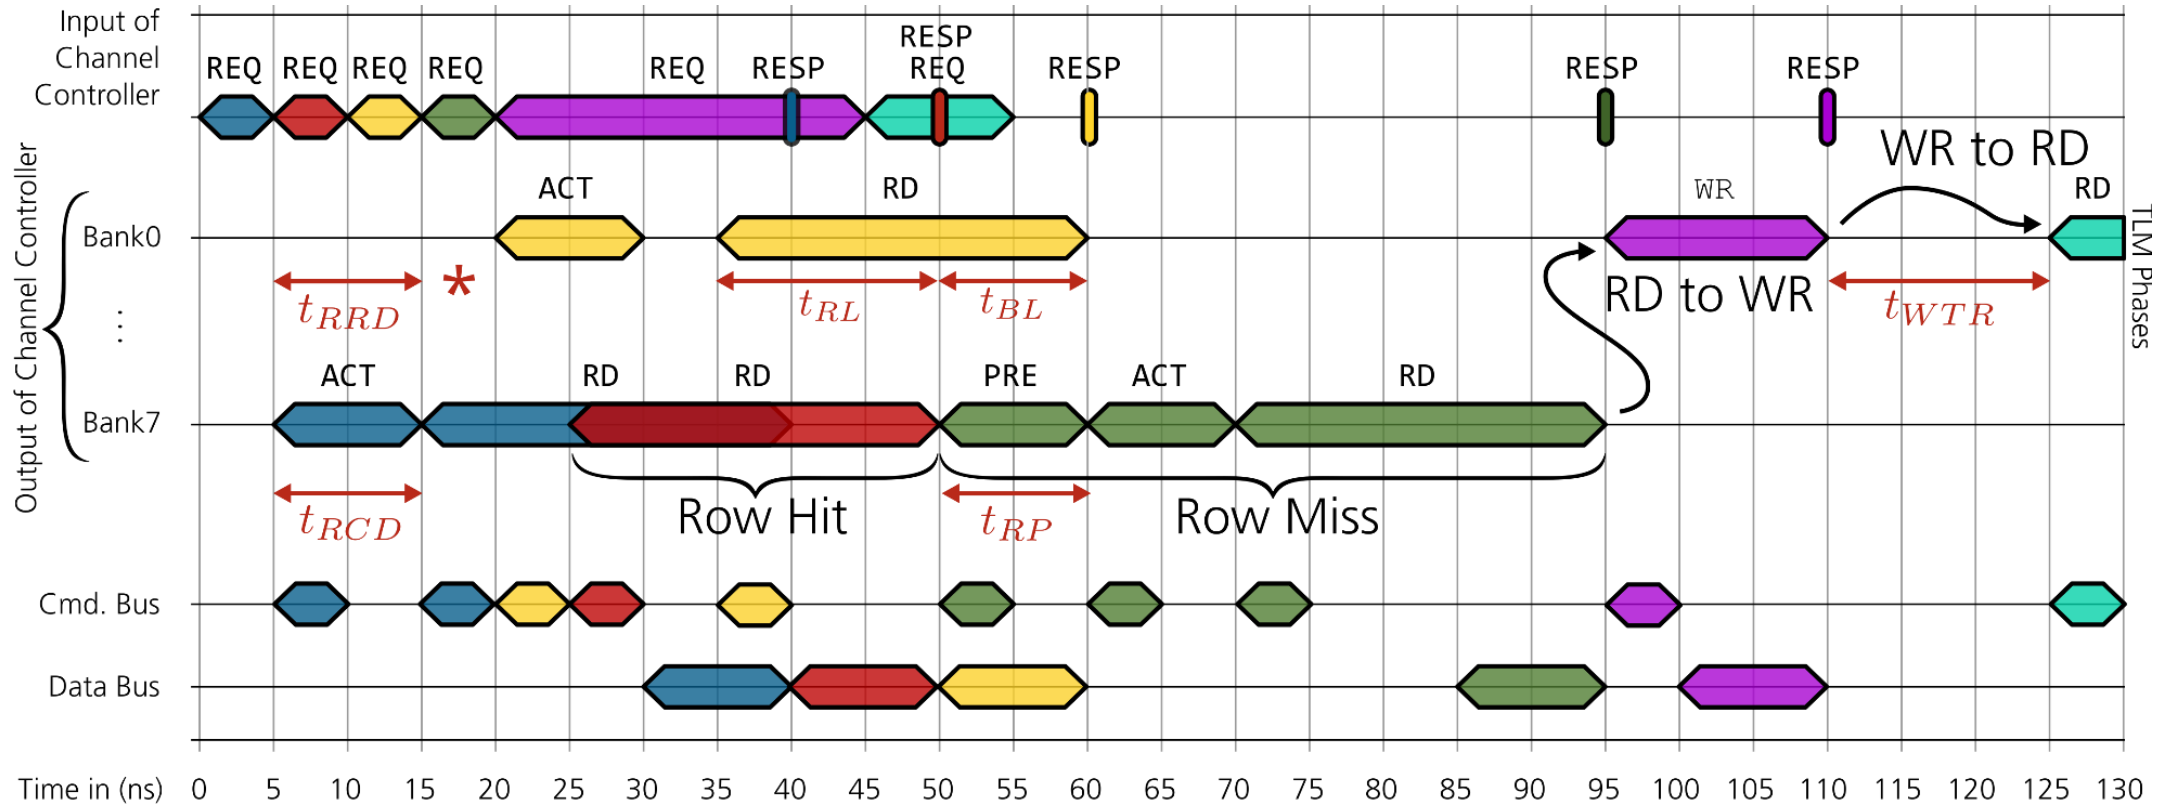
# Custom TLM Protocol

- Simulation speed can be increased by reducing the number of events
- Clock signal has the highest event generation rate
- Do we need to simulate each clock cycle to generate cycle-accurate results?

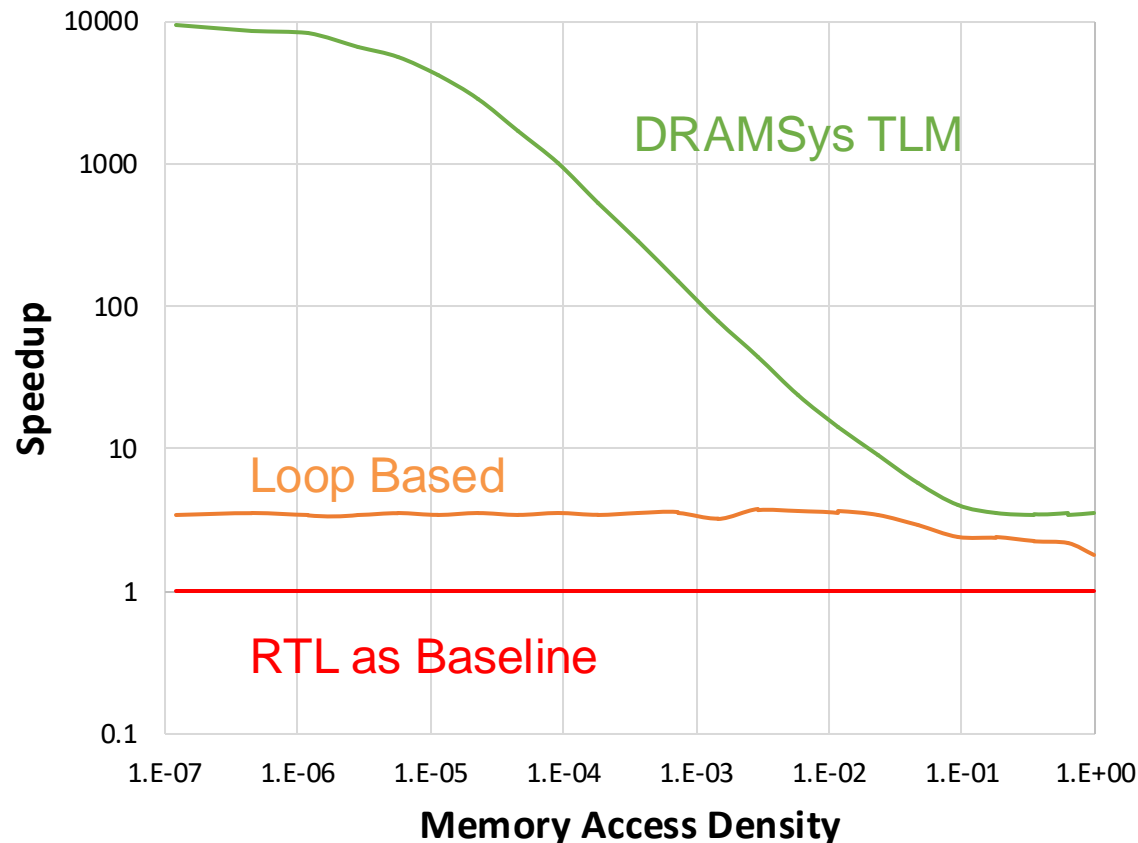


- Simulation of state changes is sufficient, idle clock cycles can be skipped!
  - Large event reduction at low memory access densities
  - No loss of accuracy

# Custom TLM Protocol

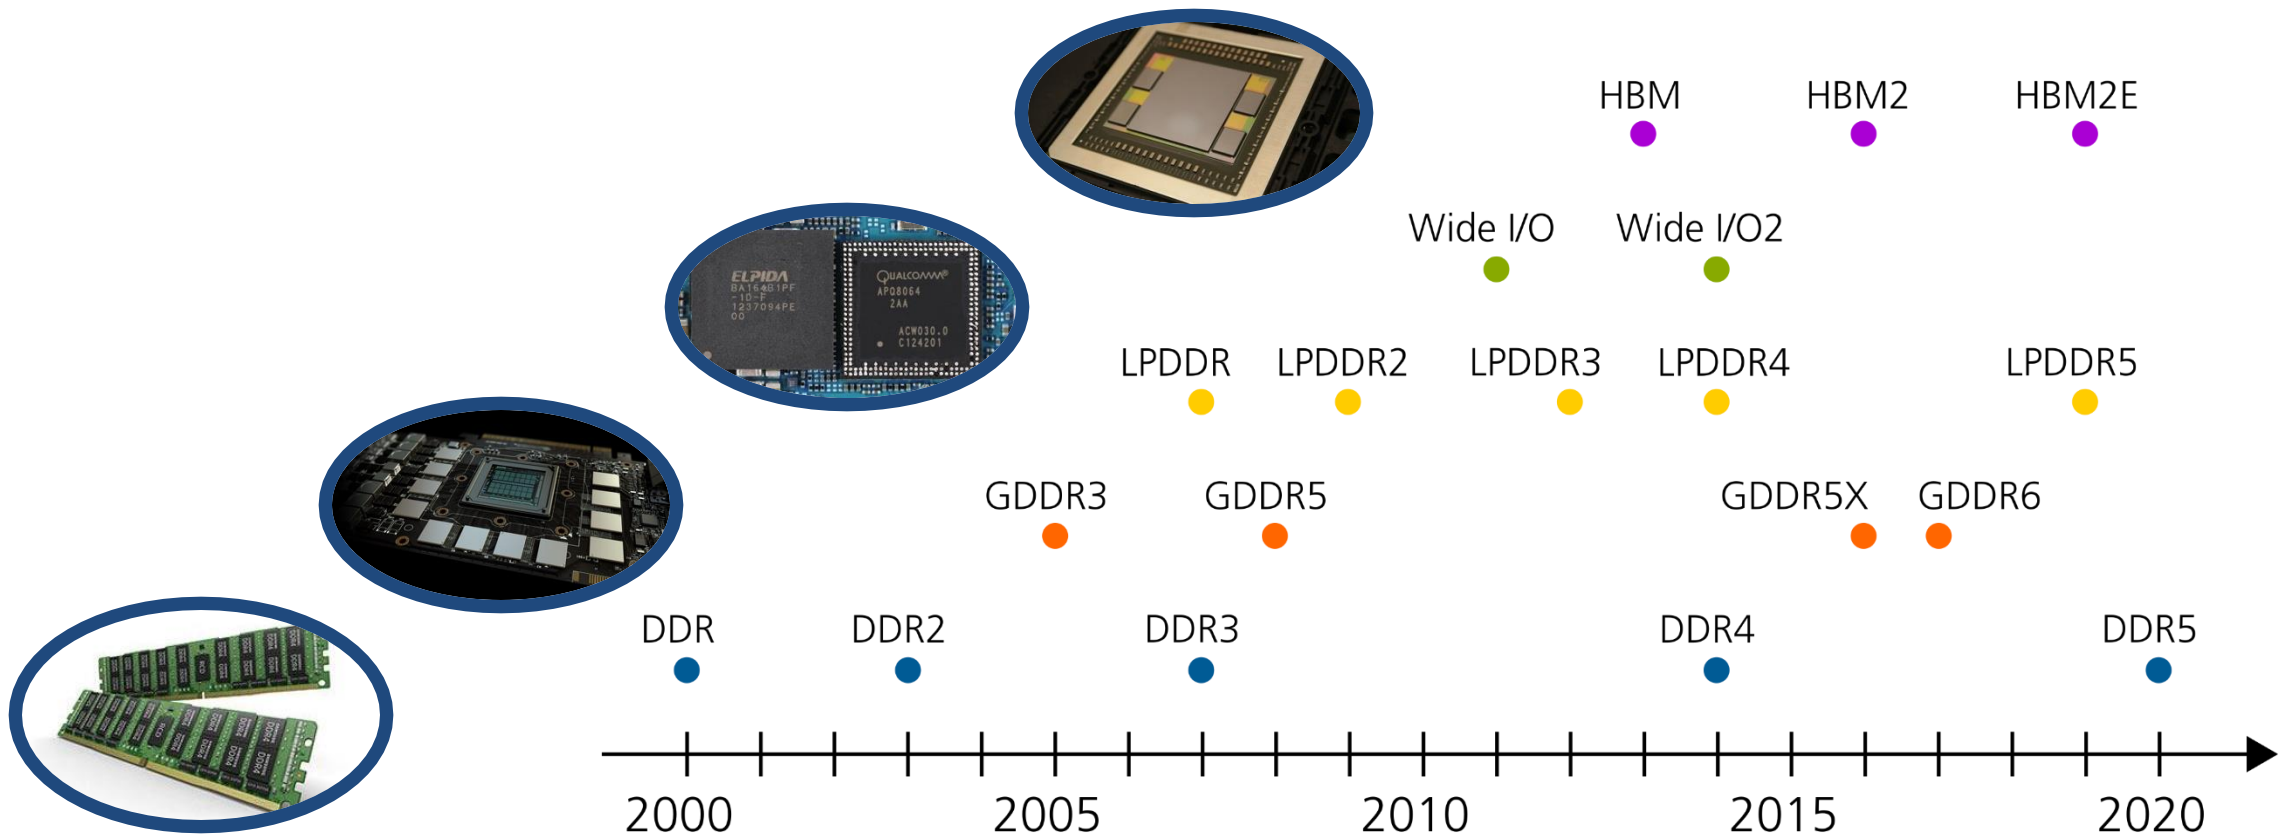


# DRAMSys Simulation Speed



- Simulation of only the important events
- Speedup from 4x to 10.000x depending on trace density
- Average speedups depend on applications
- Typical values: 400x
- 100% RTL Accuracy

# Number of DRAM Standards is Growing!

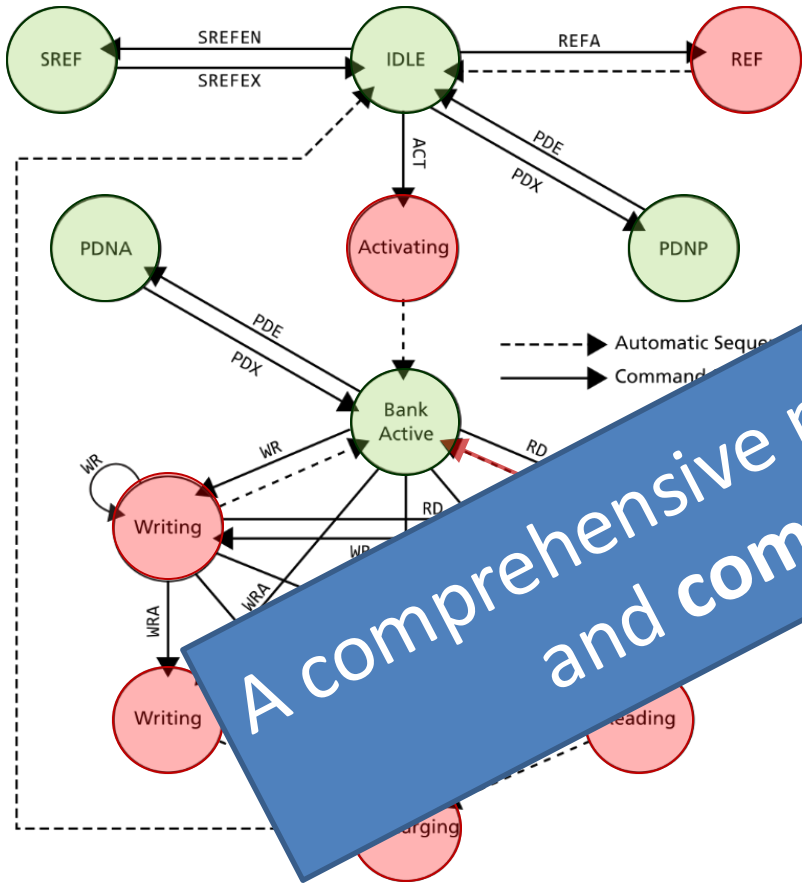


# DDR5 JEDEC Standards



- Defines commands, states, timings and interface properties
- Very complex protocol
  - DDR3: 226 pages
  - DDR4: 266 pages
  - DDR5: 496 pages
- Descriptions are not formal
- And not even correct ...

# JEDEC Standard Description “State Machine”



## DDR3 JEDEC State Machine

“This simplified description is a high-level overview of the possible state transitions and the commands used to initiate these transitions. It does not capture the timing requirements for these transitions involving more than one bank, the enabling or disabling of individual banks, or the timing of individual bank operations. These details are not captured in full detail.”


States like *Activating*, *Precharging*, *REF* ... do not exist (There are only 5 state types!)

- Double States (2x *Reading* and *Writing*)
- Inconsistencies using automatic sequences (eg. *Reading* state)



# DRAMml: a formal Description for JEDEC Standards

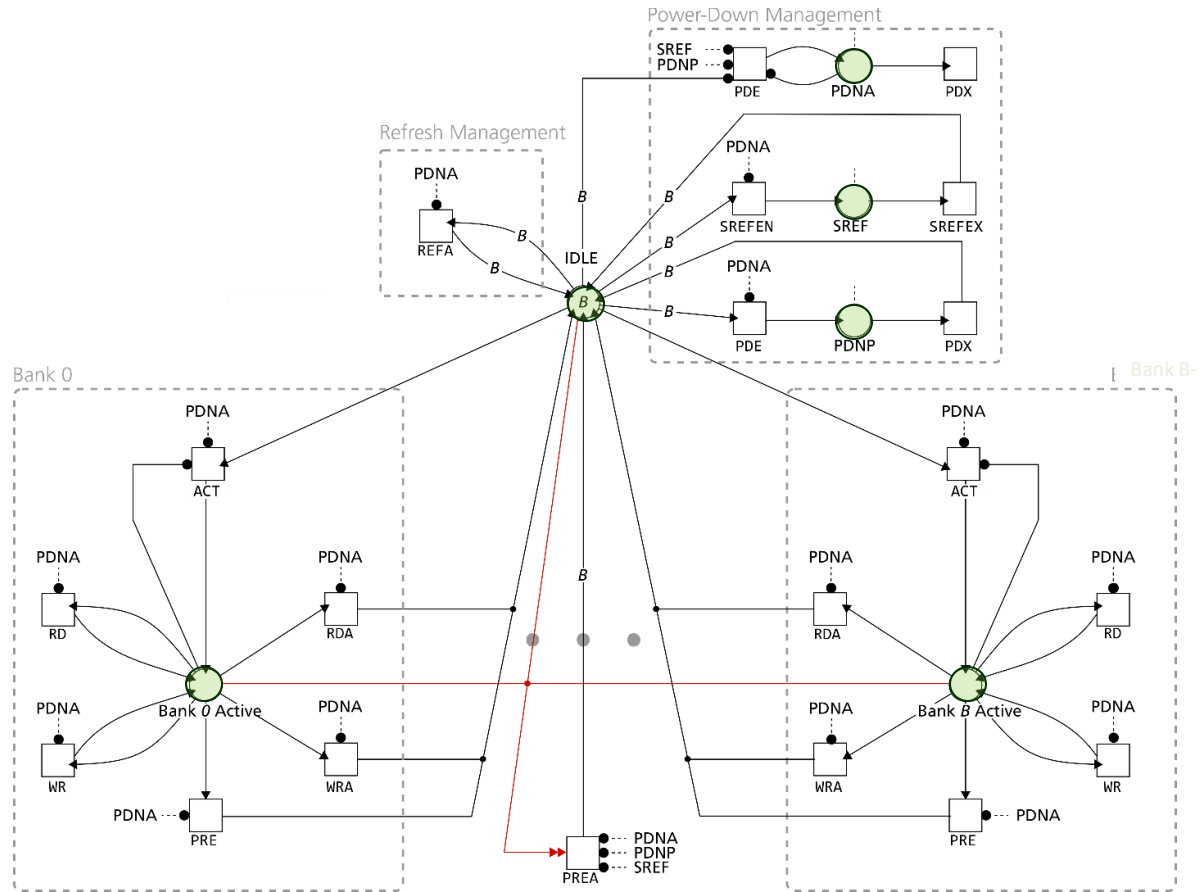
The ideal case: A *formal* language, which has the power to ...



A new standard requires a serious amount of handcraft:


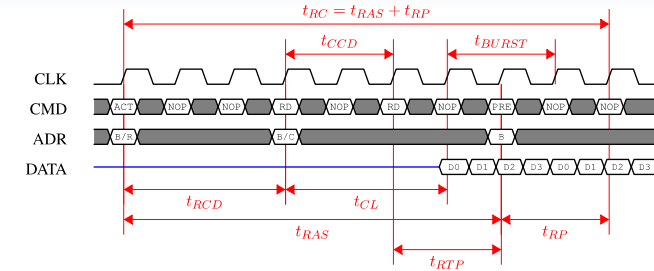
- New models for fast simulation and verification
- Adapt memory models and HW IP every time

# Modeling All DRAM States and Commands




- Comprehensive Model with clear separation between states and commands
- Models only 5 state types
- Support of multiple banks i.e. bank parallelism
- Divided in several subnets:
  - Banks
  - Refresh
  - Power-Down
  - Bankgroups
  - Ranks

# Modeling DRAM Timing

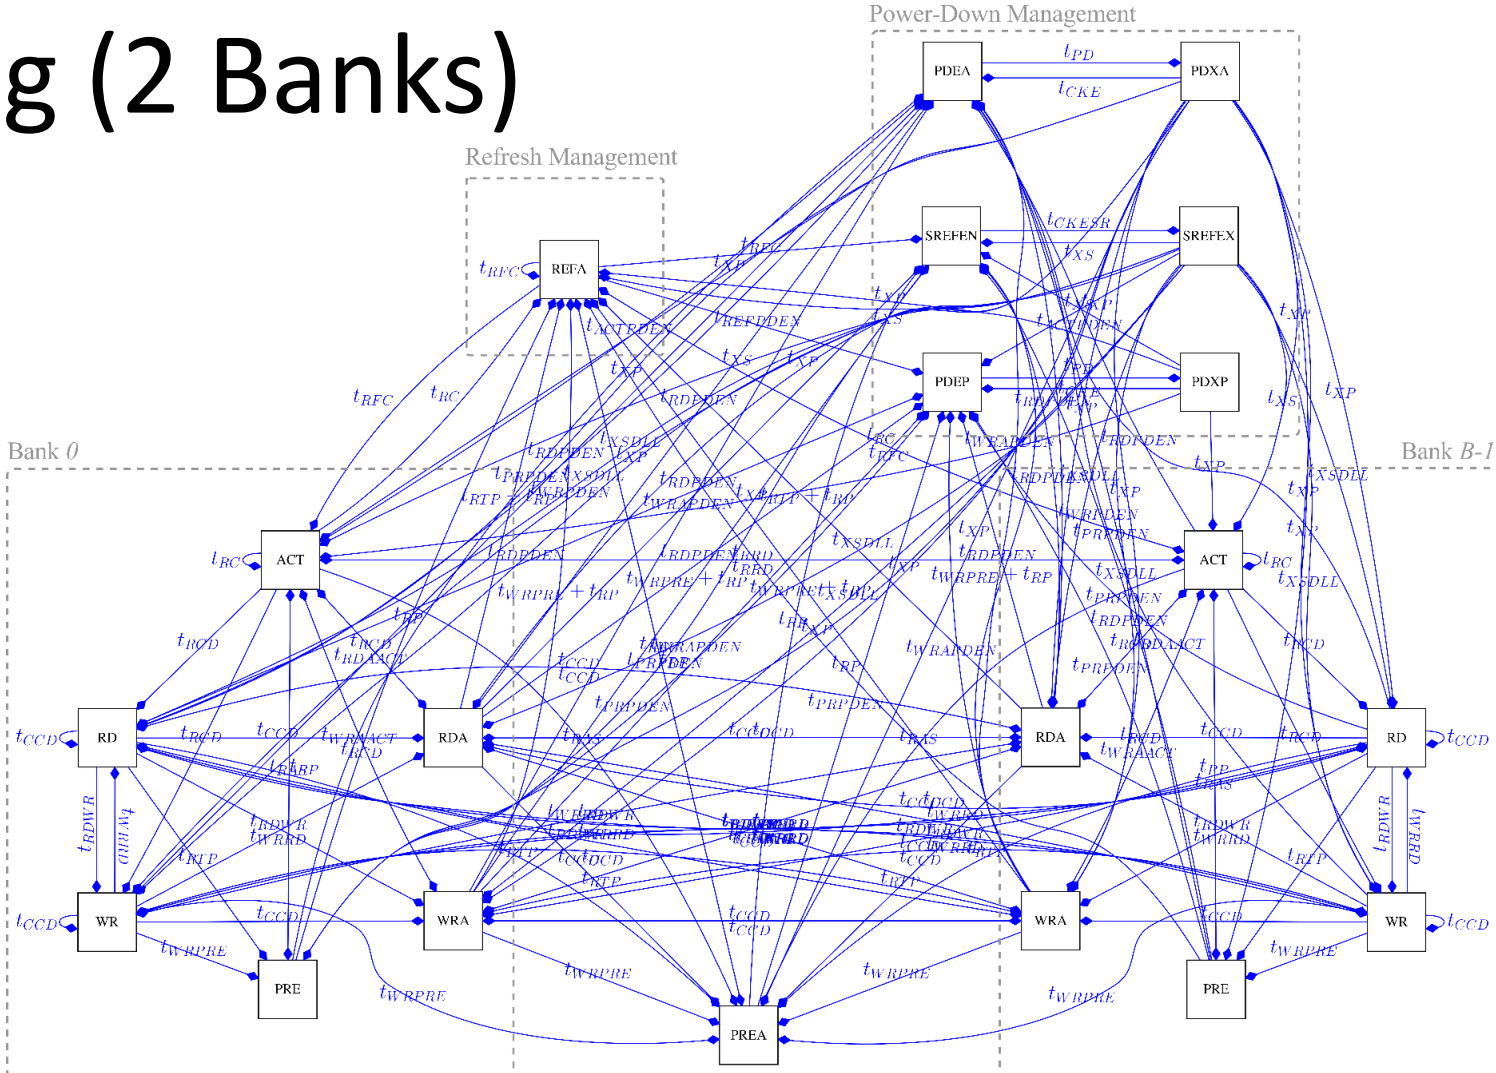


|||




- DRAMs feature a complex timing protocol
- E.g. 90 out of 260 pages of DDR3 standard are showing timing diagrams and explanations for the timings.
- DRAM command timing dependencies can be modeled by a timed inhibitor arc:
- For example, ACT to RD would be  $t_{RCD}$

# Modeling Timing (2 Banks)






# Modeling DRAM Timing (8 Banks)



# Code Generation and Validation

- Timed Petri nets allow a formal representation of a DRAM protocol, however, no graphical handling possible
- DRAMml is a DSL to describe DRAM's behavior with a petri net semantic
- DSL is as a basis for correct-by-construction DRAMSys TLM code generation
- For example: from DDR5 release it took 2 weeks to implement the model





# DDR3

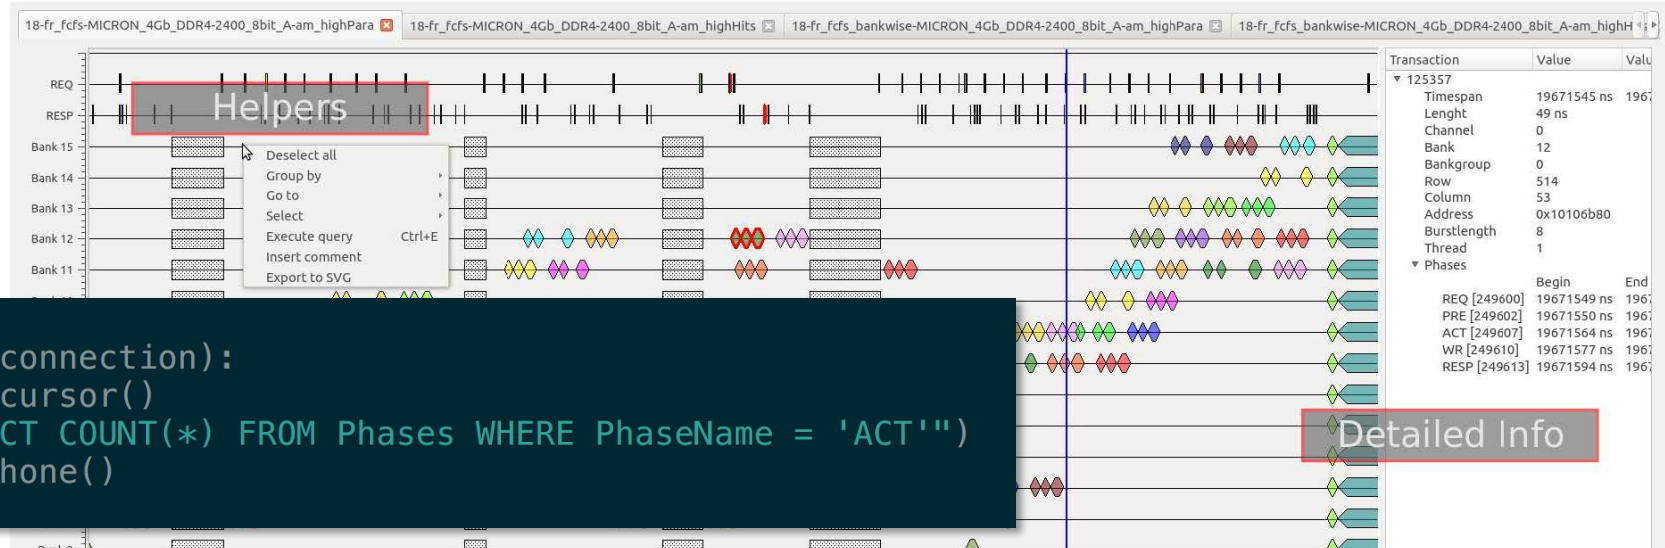
```
ddr3.py 4.81 KIB
1 from drampyl.component_levels import *
2 from drampyl.commands import *
3 from drampyl.dram import Dram
4 from drampyl.constants import *
5 from drampyl.syntax import Max
6 from drampyl.busses import CommandBus
7 from drampyl.conditions import *
8 from drampyl.constraints import *
9
10 tBURST = Variable("tBURST", defaultBurstLength / dataRate * tCK)
11 tRDWR = Variable("tRDWR", tRL + tBURST + tCK * 2 - tWL)
12 tRDWR_R = Variable("tRDWR_R", tRL + tBURST + tRTRS - tWL)
13 tWRRD = Variable("tWRRD", tWL + tBURST + tWTR - tAL)
14 tWRRD_R = Variable("tWRRD_R", tWL + tBURST + tRTRS - tRL)
15 tWRPRE = Variable("tWRPRE", tWL + tBURST + tWR)
16 tRDPDEN = Variable("tRDPDEN", tRL + tBURST + tCK)
17 tWRPDEN = Variable("tWRPDEN", tWL + tBURST + tWR)
18 tWRAPDEN = Variable("tWRAPDEN", tWL + tBURST + tWR + tCK)
19
20 custom_timings = [
21     tBURST,
22     tRDWR,
23     tRDWR_R,
24     tWRRD,
25     tWRRD_R,
26     tWRPRE,
27     tRDPDEN,
28     tWRPDEN,
29     tWRAPDEN,
30 ]
31
32 # fmt: off
33 command_timing_constraints = [
34     # Bank
35     CommandTimingConstraint(Bank, [ACT], [PREPB], tRAS),
36     CommandTimingConstraint(Bank, [ACT], [RD, WR, MWR, RDA, WRA, MWRA], tRCD - tAL),
37     CommandTimingConstraint(Bank, [ACT], [ACT], tRC),
38     CommandTimingConstraint(Bank, [RD], [PREPB], tAL + tRTP),
39     CommandTimingConstraint(Bank, [RD], [WR, MWR, WRA, MWRA], tRDWR),
40     CommandTimingConstraint(Bank, [RDA], [ACT], tAL + tRTP + tRP),
41     CommandTimingConstraint(Bank, [WR, MWR], [PREPB], tWRPRE),
42     CommandTimingConstraint(Bank, [WR, MWR], [WR, MWR, WRA, MWRA], tCCD),
43     CommandTimingConstraint(Bank, [WR, MWR], [RD], tWRRD),
44     CommandTimingConstraint(Bank, [WR, MWR], [RDA], Max(tWRRD, tWRPRE - tRTP - tAL)),
45     CommandTimingConstraint(Bank, [WRA, MWRA], [ACT], tWRPRE + tRP),
46     CommandTimingConstraint(Bank, [PREPB], [ACT], tRP),
47
48     # Rank
49     CommandTimingConstraint(Rank, [ACT], [PREAB], tRAS),
50     CommandTimingConstraint(Rank, [ACT], [ACT], tRRD),
51     CommandTimingConstraint(Rank, [ACT], [PDEA], tACTPDEN),
52     CommandTimingConstraint(Rank, [ACT], [REFAB, SREFEN], tRC),
53     CommandTimingConstraint(Rank, [RD], [PREAB], tAL + tRTP),
54     CommandTimingConstraint(Rank, [RD, RDA], [PDEA, PDEP], tRDPDEN),
55     CommandTimingConstraint(Rank, [RD, RDA], [RD, RDA], tCCD),
56     CommandTimingConstraint(Rank, [RD, RDA], [WR, MWR, WRA, MWRA], tRDWR),
57     CommandTimingConstraint(Rank, [RDA], [REFAB], tAL + tRTP + tRP),
58     CommandTimingConstraint(Rank, [RDA], [PREAB], tAL + tRTP),
59     CommandTimingConstraint(Rank, [RDA], [SREFEN], Max(tRDPDEN, tAL + tRTP + tRP)),
60     CommandTimingConstraint(Rank, [WR, MWR], [PDEA], tWRPDEN),
61     CommandTimingConstraint(Rank, [WRA, MWRA], [PDEA, PDEP], tWRAPDEN),
62     CommandTimingConstraint(Rank, [WR, MWR, WRA, MWRA], [WR, MWR, WRA, MWRA], tCCD),
63     CommandTimingConstraint(Rank, [WR, MWR, WRA, MWRA], [RD, RDA], tWRRD),
64     CommandTimingConstraint(Rank, [WRA, MWRA], [REFAB], tWRPRE + tRP),
65     CommandTimingConstraint(Rank, [WRA, MWRA], [PREAB], tWRPRE),
66     CommandTimingConstraint(Rank, [WRA, MWRA], [SREFEN], Max(tWRAPDEN, tWRPRE + tRP)),
67     CommandTimingConstraint(Rank, [PREPB], [REFAB], tRP),
68     CommandTimingConstraint(Rank, [PREPB], [PDEA, PDEP], tPRPDEN),
69     CommandTimingConstraint(Rank, [PREPB], [SREFEN], tRP),
70     CommandTimingConstraint(Rank, [PREAB], [ACT, REFAB, SREFEN], tRP),
71     CommandTimingConstraint(Rank, [PREAB], [PDEP], tPRPDEN),
72     CommandTimingConstraint(Rank, [PDEP], [PDXP], tPD),
73     CommandTimingConstraint(Rank, [PDEA], [PDXA], tPD),
74     CommandTimingConstraint(Rank, [PDXA], [PDEA], tCKE),
75     CommandTimingConstraint(Rank, [PDXA], [PDEP], tCKE),
76     CommandTimingConstraint(Rank, [PDXP], [REFAB, SREFEN, ACT], tXP),
77     CommandTimingConstraint(Rank, [PDXA], [ACT, PREPB, PREAB, RD, RDA, WR, MWR, WRA, MWRA], tXP),
78     CommandTimingConstraint(Rank, [REFAB], [ACT, REFAB, SREFEN], tRFC),
79     CommandTimingConstraint(Rank, [REFAB], [PDEP], tREFPDEN),
80     CommandTimingConstraint(Rank, [SREFEX], [ACT, REFAB, PDEP, SREFEN], tXS),
81     CommandTimingConstraint(Rank, [SREFEX], [RD, RDA, WR, MWR, WRA, MWRA], tXSDLL),
82     CommandTimingConstraint(Rank, [SREFEX], [SREFEX], tCKESR),
83
84     # Channel
85     CommandTimingConstraint(Channel, [RD, RDA], [RD, RDA], tBURST + tRTRS, [], Different(Rank)),
86     CommandTimingConstraint(Channel, [RD, RDA], [WR, MWR, WRA, MWRA], tRDWR_R, [], Different(Rank)),
87     CommandTimingConstraint(Channel, [WR, MWR, WRA, MWRA], [WR, MWR, WRA, MWRA], tBURST + tRTRS, [], Different(Rank)),
88     CommandTimingConstraint(Channel, [WR, MWR, WRA, MWRA], [RD, RDA], tWRRD_R, [], Different(Rank)),
89 ]
90 # fmt: on
91
92 faw_constraints = [FANConstraint([ACT], Rank, tFAW)]
93
94 bus_commands = {
95     ACT,
96     RD,
97     RDA,
98     WR,
99     WRA,
100     MWR,
101     MWRA,
102     PREPB,
103     PREAB,
104     PDEP,
105     PDEA,
106     PDXA,
107     PDXP,
108     REFAB,
109     SREFEX,
110     SREFEN,
111 }
112
113 command_bus = CommandBus("Bus", bus_commands)
114
115 dram = Dram(
116     "DDR3",
117     [command_bus],
118     command_timing_constraints,
119     faw_constraints,
120     custom_timings,
121 )
122
```

```

62 CommandTimingConstraint(Rank, [WR, MWR, WRA, MWRA], [WR, MWR, WRA, MWRA], tCCD),
63 CommandTimingConstraint(Rank, [WR, MWR, WRA, MWRA], [RD, RDA], tWRRD),
64 CommandTimingConstraint(Rank, [WRA, MWRA], [REFAB], tWRPRE + tRP),
65 CommandTimingConstraint(Rank, [WRA, MWRA], [PREAB], tWRPRE),
66 CommandTimingConstraint(Rank, [WRA, MWRA], [SREFEN], Max(tWRAPDEN, tWRPRE + tRP)),
67 CommandTimingConstraint(Rank, [PREPB], [REFAB], tRP),
68 CommandTimingConstraint(Rank, [PREPB], [PDEA, PDEP], tPRPDEN),
69 CommandTimingConstraint(Rank, [PREPB], [SREFEN], tRP),
70 CommandTimingConstraint(Rank, [PREAB], [ACT, REFAB, SREFEN], tRP),
71 CommandTimingConstraint(Rank, [PREAB], [PDEP], tPRPDEN),
72 CommandTimingConstraint(Rank, [PDEP], [PDXP], tPD),
73 CommandTimingConstraint(Rank, [PDEA], [PDXA], tPD),
74 CommandTimingConstraint(Rank, [PDXA], [PDEA], tCKE),
75 CommandTimingConstraint(Rank, [PDXA], [PDEP], tCKE),
76 CommandTimingConstraint(Rank, [PDXP], [REFAB, SREFEN, ACT], tXP),
77 CommandTimingConstraint(Rank, [PDXA], [ACT, PREPB, PREAB, RD, RDA, WR, MWR, WRA, MWRA], tXP),
78 CommandTimingConstraint(Rank, [REFAB], [ACT, REFAB, SREFEN], tRFC),
79 CommandTimingConstraint(Rank, [REFAB], [PDEP], tREFPDEN),
80 CommandTimingConstraint(Rank, [SREFEX], [ACT, REFAB, PDEP, SREFEN], tXS),
81 CommandTimingConstraint(Rank, [SREFEX], [RD, RDA, WR, MWR, WRA, MWRA], tXSDLL),
82 CommandTimingConstraint(Rank, [SREFEX], [SREFEX], tCKESR),
83
84 # Channel
85 CommandTimingConstraint(Channel, [RD, RDA], [RD, RDA], tBURST + tRTRS, [], Different(Rank)),
86 CommandTimingConstraint(Channel, [RD, RDA], [WR, MWR, WRA, MWRA], tRDWR_R, [], Different(Rank)),
87 CommandTimingConstraint(Channel, [WR, MWR, WRA, MWRA], [WR, MWR, WRA, MWRA], tBURST + tRTRS, [], Different(Rank)),
88 CommandTimingConstraint(Channel, [WR, MWR, WRA, MWRA], [RD, RDA], tWRRD_R, [], Different(Rank)),
89 ]
90 # fmt: on
91
92 faw_constraints = [FANConstraint([ACT], Rank, tFAW)]
93
94 bus_commands = {
95     ACT,
96     RD,
97     RDA,
98     WR,
99     WRA,
100     MWR,
101     MWRA,
102     PREPB,
103     PREAB,
104     PDEP,
105     PDEA,
106     PDXA,
107     PDXP,
108     REFAB,
109     SREFEX,
110     SREFEN,
111 }
112
113 command_bus = CommandBus("Bus", bus_commands)
114
115 dram = Dram(
116     "DDR3",
117     [command_bus],
118     command_timing_constraints,
119     faw_constraints,
120     custom_timings,
121 )
122
```



# Trace Analyzer



```

215 @metric
216 def number_of_activates(connection):
217     cursor = connection.cursor()
218     cursor.execute("SELECT COUNT(*) FROM Phases WHERE PhaseName = 'ACT'")
219     result = cursor.fetchone()
220     return result[0]
    
```

**Evaluate Traces**

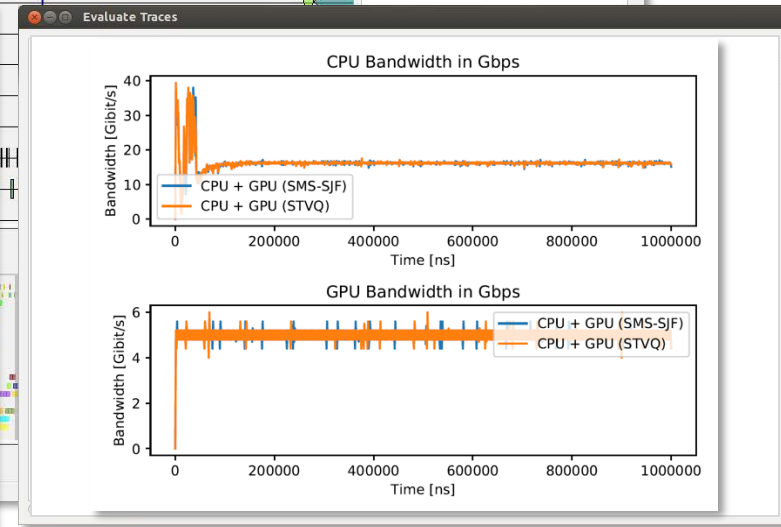
Test results for various configurations:

- 13-fr\_fcfs-MICRON\_4Gb\_DDR4-2400\_8bit\_A-am\_highPara: All tests passed!
- 13-fr\_fcfs\_bankwise-MICRON\_4Gb\_DDR4-2400\_8bit\_A-am\_highPara: All tests passed!
- 14-fr\_fcfs-MICRON\_4Gb\_DDR4-2400\_8bit\_A-am\_highPara: All tests passed!
- 14-fr\_fcfs\_bankwise-MICRON\_4Gb\_DDR4-2400\_8bit\_A-am\_highPara: All tests passed!

**Evaluate Traces**

Calculate Metrics

- 13-fr\_fcfs-MICRON\_4Gb\_DDR4-2400\_8bit\_A-am\_highPara**
  - average response latency in ns: 84.3
  - median response latency in ns: 37.6
  - number of activates: 16454
  - accesses per activate: 2.1
  - Active time (%): 6.79003
  - Time in PDNA (%): 27.334
  - Time in PDNP (%): 23.5452
  - Time in SREF (%): 42.3307
- 13-fr\_fcfs\_bankwise-MICRON\_4Gb\_DDR4-2400\_8bit\_A-am\_highPara**
  - average response latency in ns: 126.1
  - median response latency in ns: 37.5
  - number of activates: 15623
  - number of precharges: 15614
  - accesses per activate: 2.2
  - Active time (%): 1.00676
  - Time in PDNA (%): 4.56337
  - Time in PDNP (%): 3.38575
  - Time in SREF (%): 91.0441
- 14-fr\_fcfs-MICRON\_4Gb\_DDR4-2400\_8bit\_A-am\_highPara**
  - average response latency in ns: 307.7
  - median response latency in ns: 304.9



# Where to get the tools:

**DRAMSys:**

<https://github.com/tukl-msd/DRAMSys>

**gem5:**

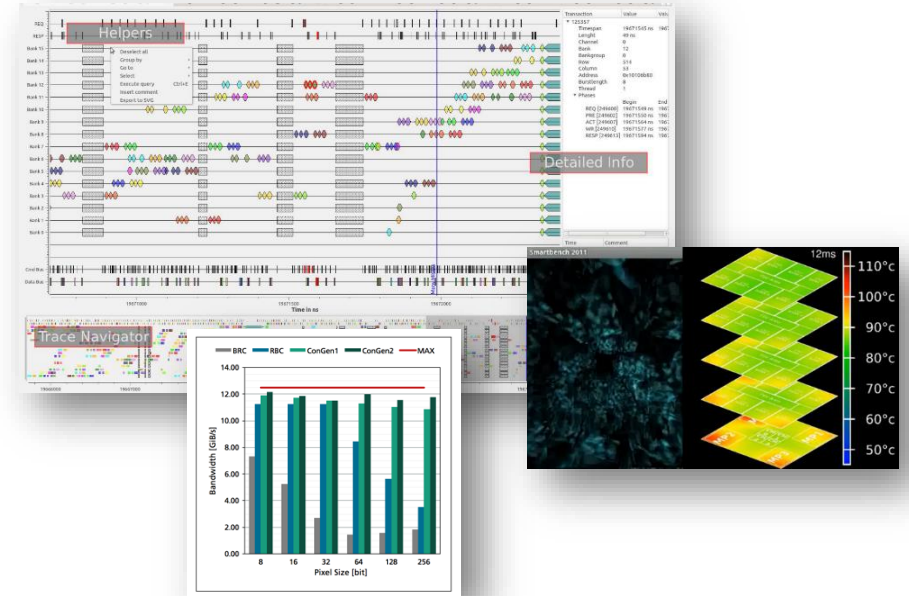
<https://gem5.googlesource.com/>

**DRAMPower:**

<https://github.com/tukl-msd/DRAMPower>

**3D-ICE:**

<https://github.com/esl-epfl/3d-ice>





## The Open Source DRAM Simulator DRAMSys

Prof. Dr. Matthias Jung, JMU Würzburg

