Safe Limits on Voltage Reduction Efficiency in GPUs: a Direct Measurement Approach

Jingwen Leng,
Alper Buyuktosunoglu, Ramon Bertran, Pradip Bose, Vijay Janapa Reddi

This work is sponsored in part by Defense Advanced Research Projects Agency (DARPA), Microsystems Technology Office (MTO), under contract number HR0011-13-C-0022, National Science Foundation (NSF), under grant CCF-1218474, and Semiconductor Research Corporation (SRC). The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense, the NSF, the SRC or the U.S. Government. This document is: Approved for Public Release, Distribution Unlimited.

GPU Energy Efficiency Optimization

GPU Energy Efficiency Optimization

Circuit –

DVFS
Clock Gating
Power Gating

GPU Energy Efficiency Optimization

(Micro)architecture

Control Divergence
Warp Scheduler
Cache Locality

Circuit

DVFS
Clock Gating
Power Gating

GPU Energy Efficiency Optimization

Software/
Compiler

Approximation
Data Transfer
Multi-tasking

(Micro)architecture

Control Divergence
Warp Scheduler
Cache Locality

Circuit

DVFS
Clock Gating
Power Gating

GPU Energy Efficiency Optimization

Software/ Compiler

Approximation
Data Transfer
Multi-tasking

Samadi et al. [MICRO'14]
Rossbach et al. [SOSP'11]
Park et al. [ASPLOS'15]

(Micro)architecture

Control Divergence
Warp Scheduler
Cache Locality

Fung et al. [MICRO'07]
Rogers et al. [MICRO'13]
Rhu et al. [MICRO'13]

Circuit

DVFS
Clock Gating
Power Gating

Sethia et al. [MICRO'14]
Leng et al. [ISCA'13]
Majeed et al. [MICRO'13]

Energy Inefficiency at the Voltage Guardband

Energy Inefficiency at the Voltage Guardband

Operating
Supply Voltage

Voltage
Guardband

Required
Supply Voltage

Energy Inefficiency at the Voltage Guardband

Operating
Supply Voltage

Voltage
Guardband

Required
Supply Voltage

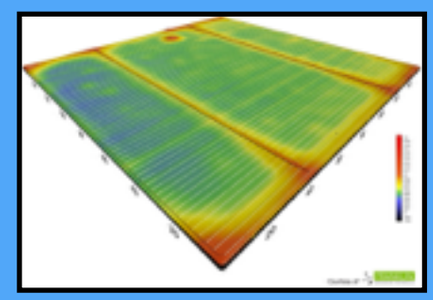

Process

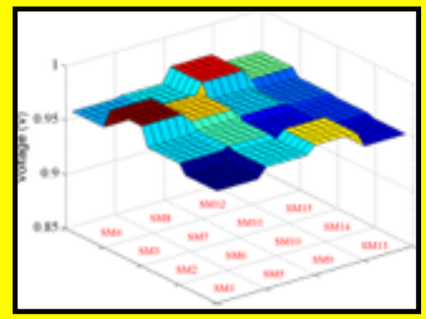

Voltage

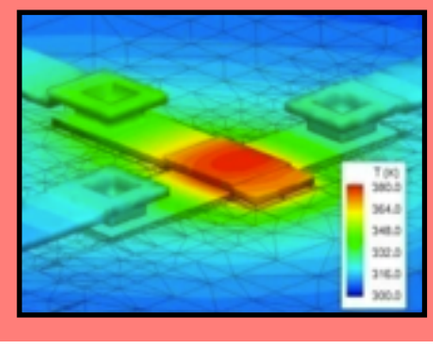

Tempera
-ture

Energy Inefficiency at the Voltage Guardband

Operating
Supply Voltage

Voltage
Guardband

Required
Supply Voltage

Energy Inefficiency at the Voltage Guardband

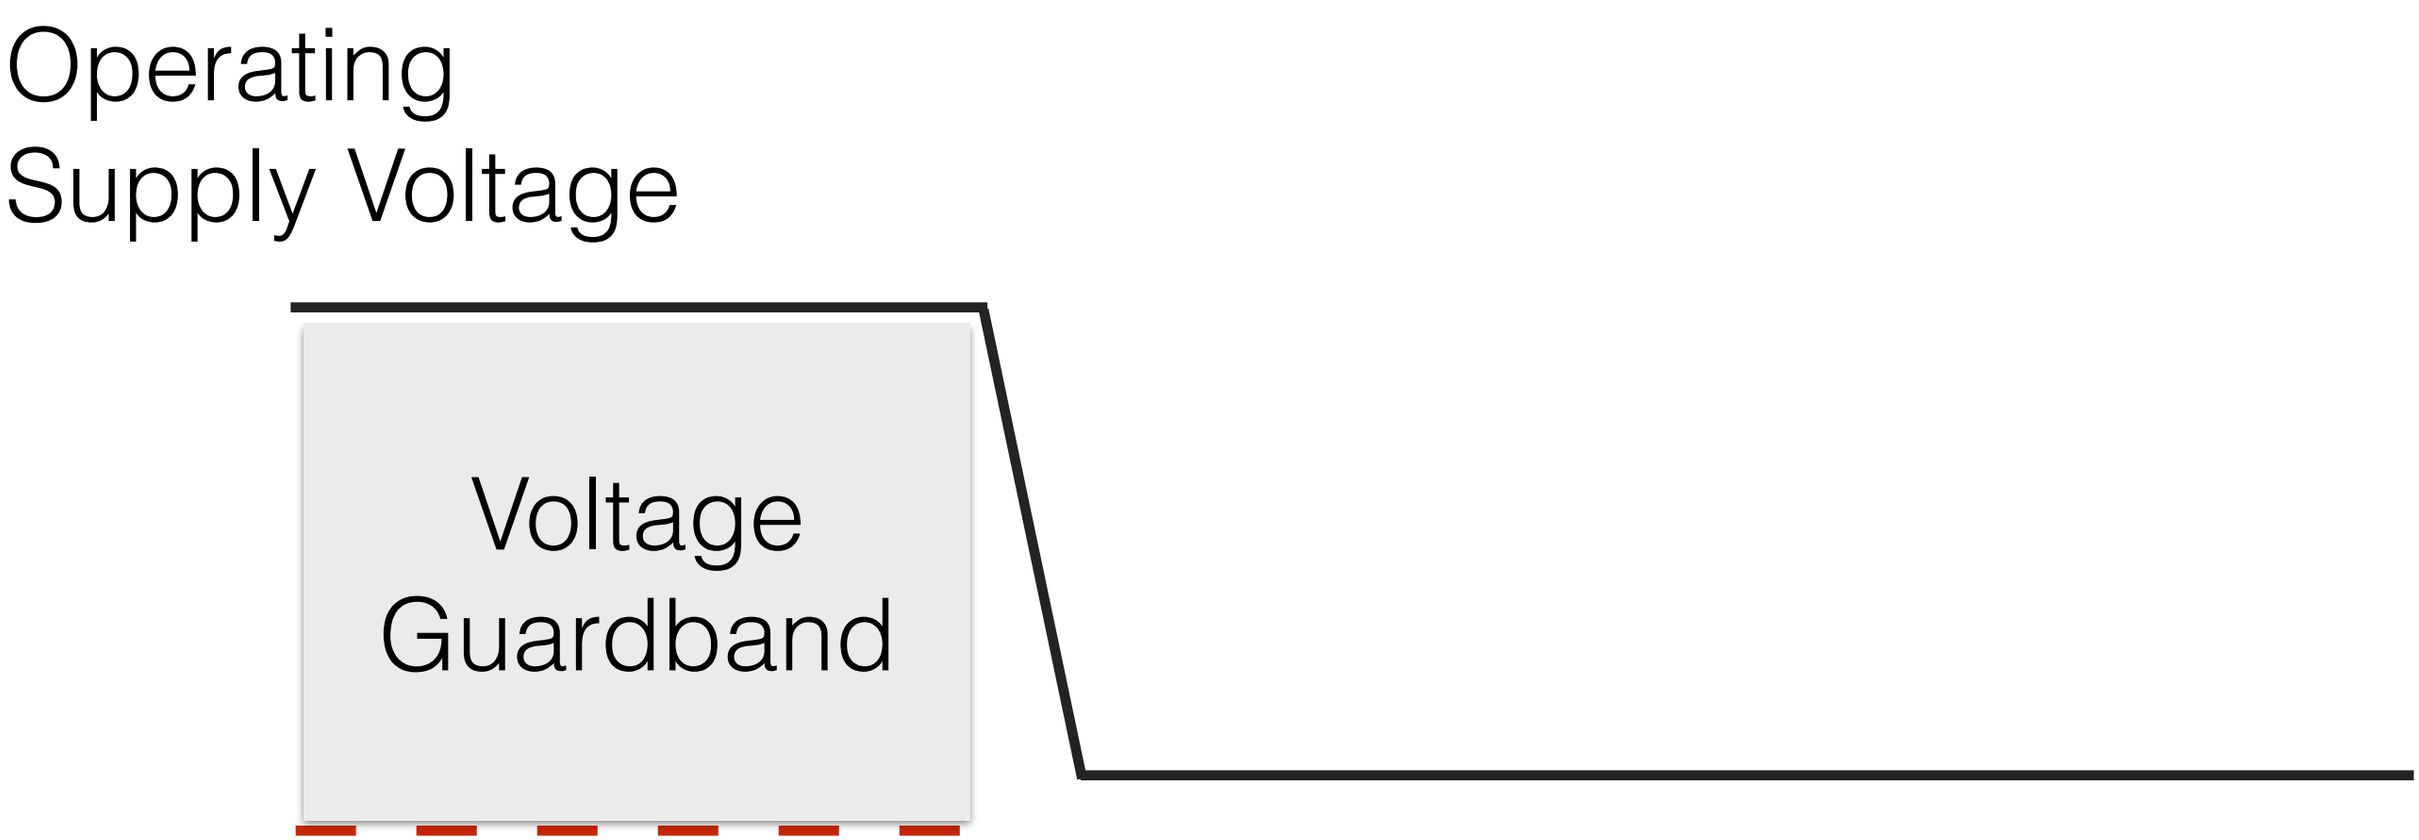

Energy Inefficiency at the Voltage Guardband

Operating
Supply Voltage

Voltage
Guardband

Reduced voltage → energy savings

Required
Supply Voltage

Our Contributions

Our Contributions

Our Contributions

- Voltage guardband measurement

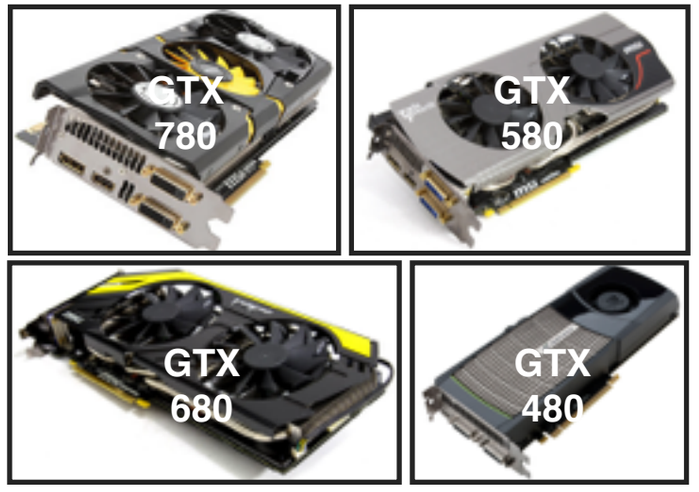

Our Contributions

- Voltage guardband measurement

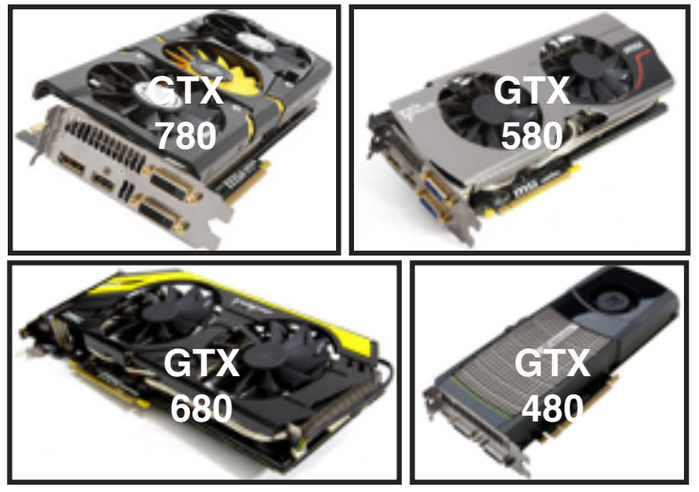

Our Contributions

- Voltage guardband measurement
- Guardband analysis

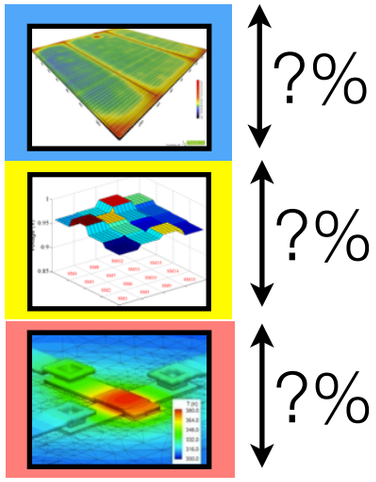

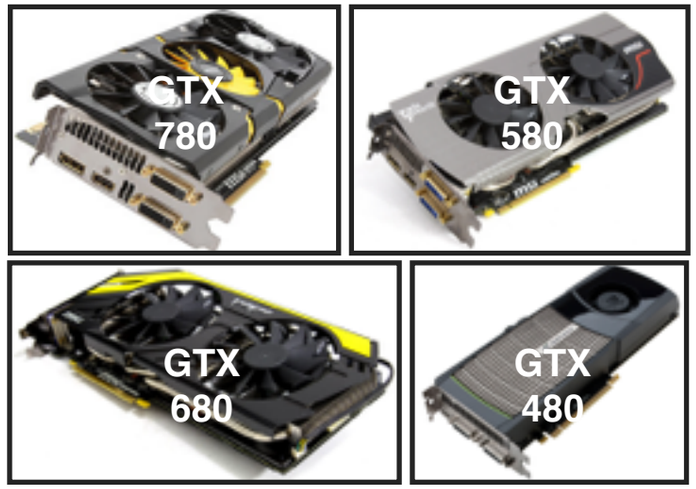

Our Contributions

- Voltage guardband measurement

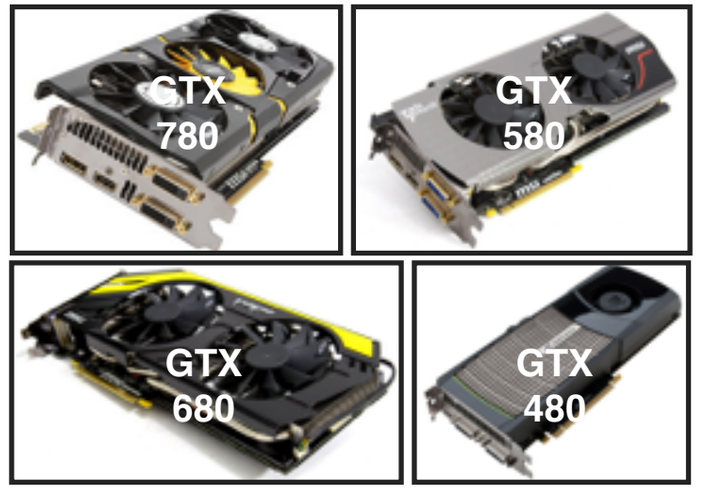

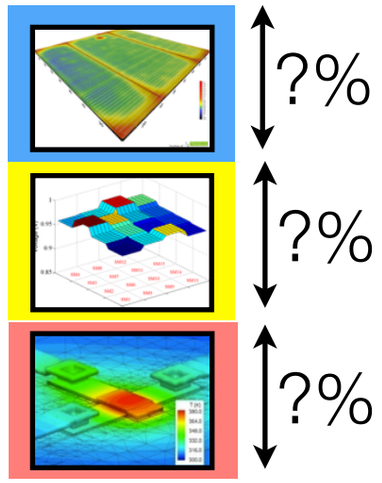

- Guardband analysis

- Program-driven predictive guardbanding

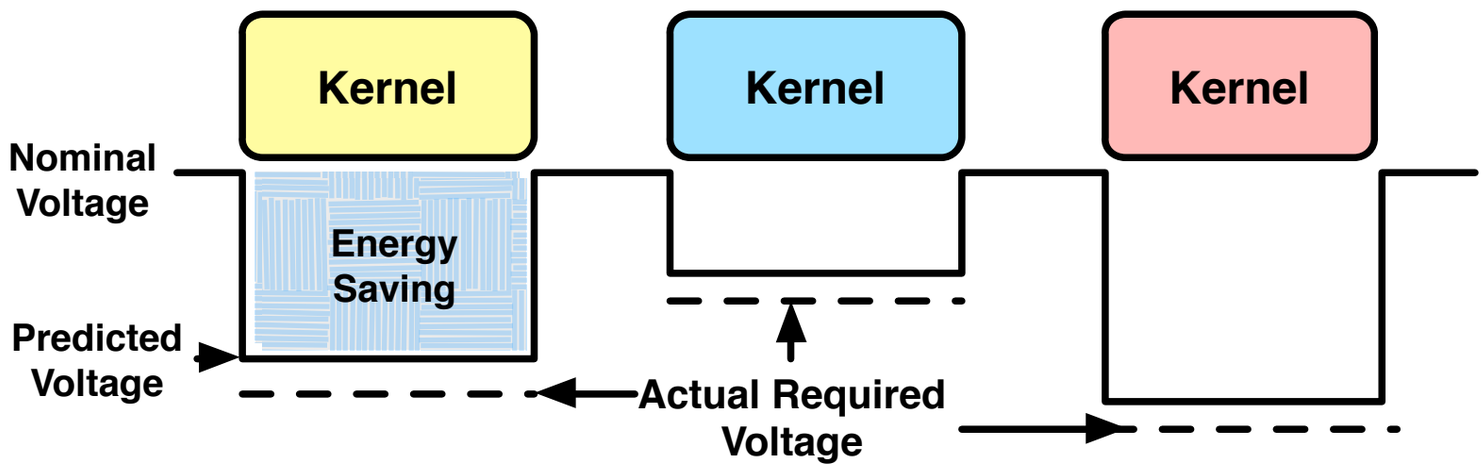

Voltage Guardband Measurement

Voltage Guardband Measurement

- Eight GPU cards in total
  - Four generations
  - Two different architectures

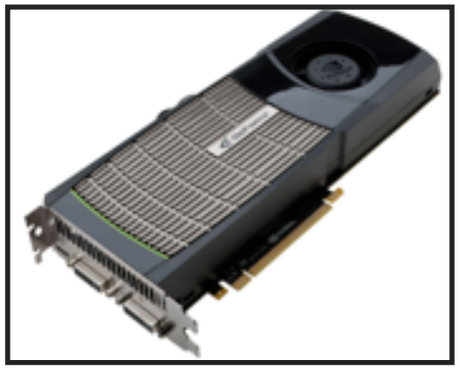

GTX 480 x1

GTX 580 x1

GTX 680 x1

GTX 780 x5

Voltage Guardband Measurement

- Eight GPU cards in total
  - Four generations
  - Two different architectures
- Fifty-seven representative CUDA programs
  - Regular/irregular
  - Memory/arithmetic intensive

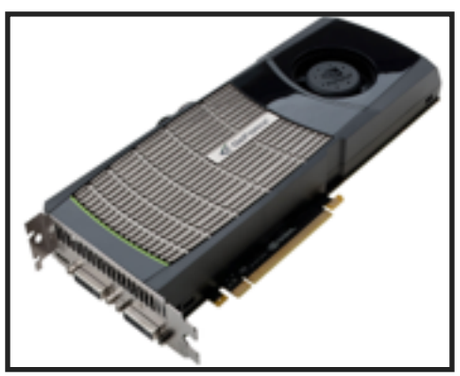

GTX 480 x1

GTX 580 x1

GTX 680 x1

GTX 780 x5

V_{min} Measurement

CUDA Programs

V_{min} Measurement

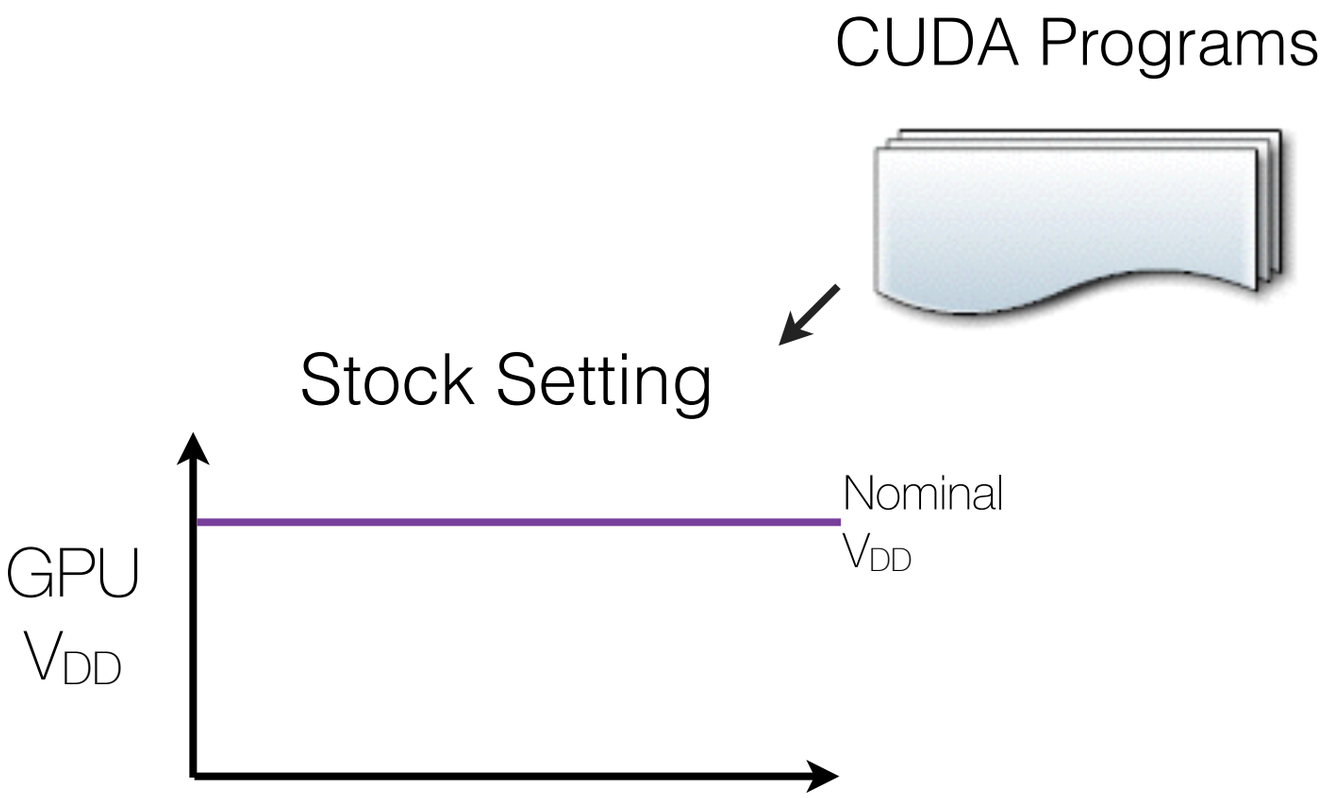

V_{min} Measurement

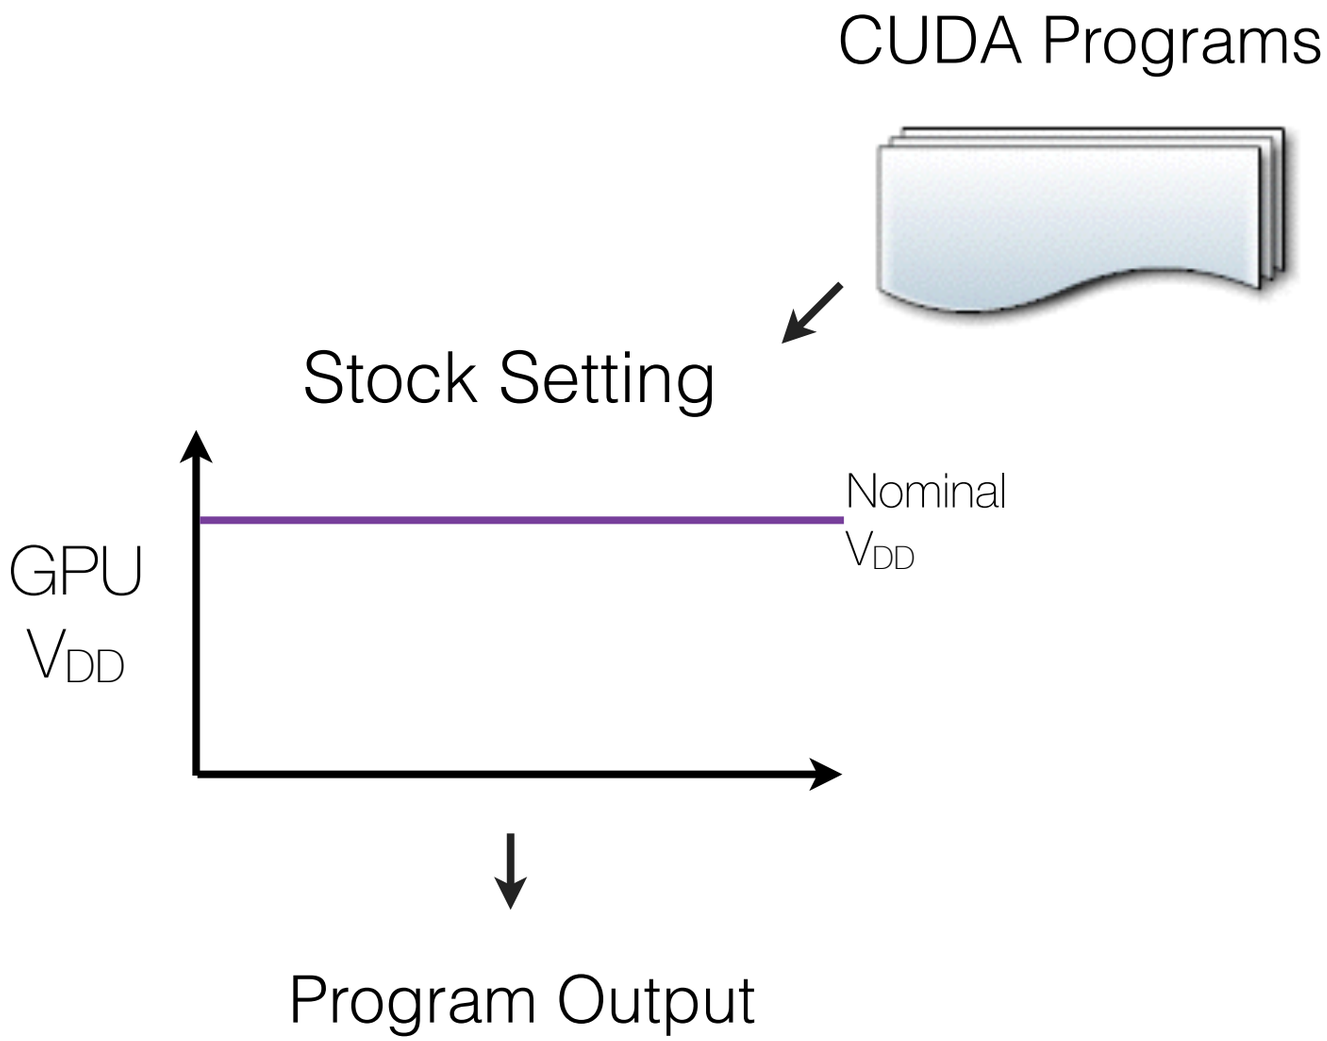

V_{min} Measurement

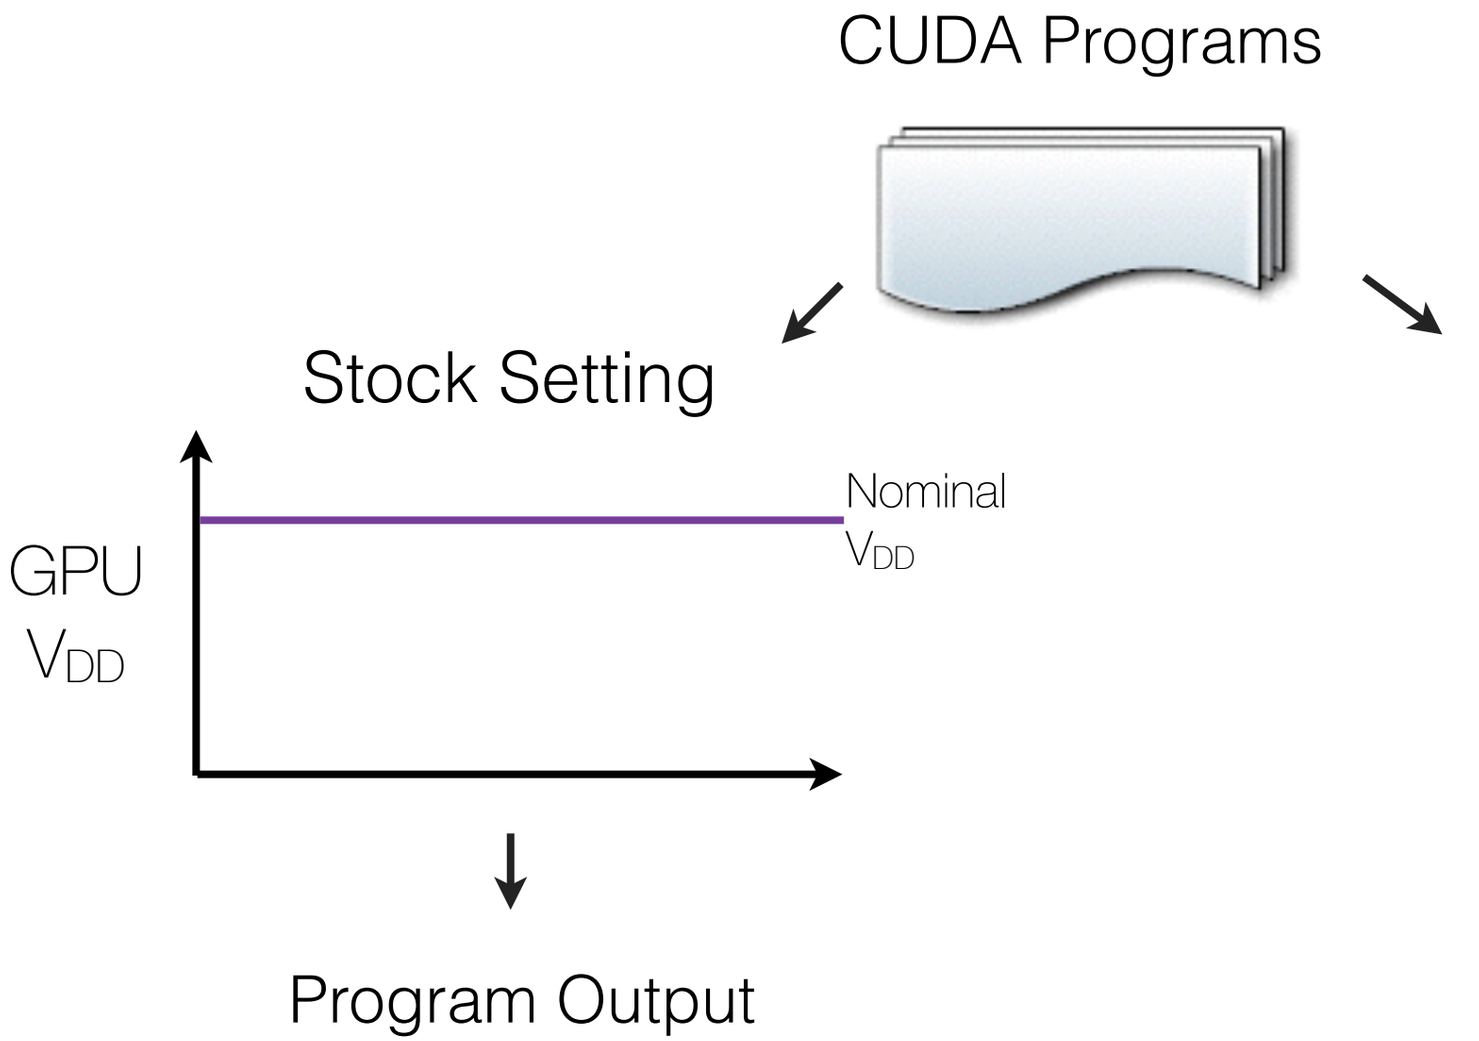

V_{min} Measurement

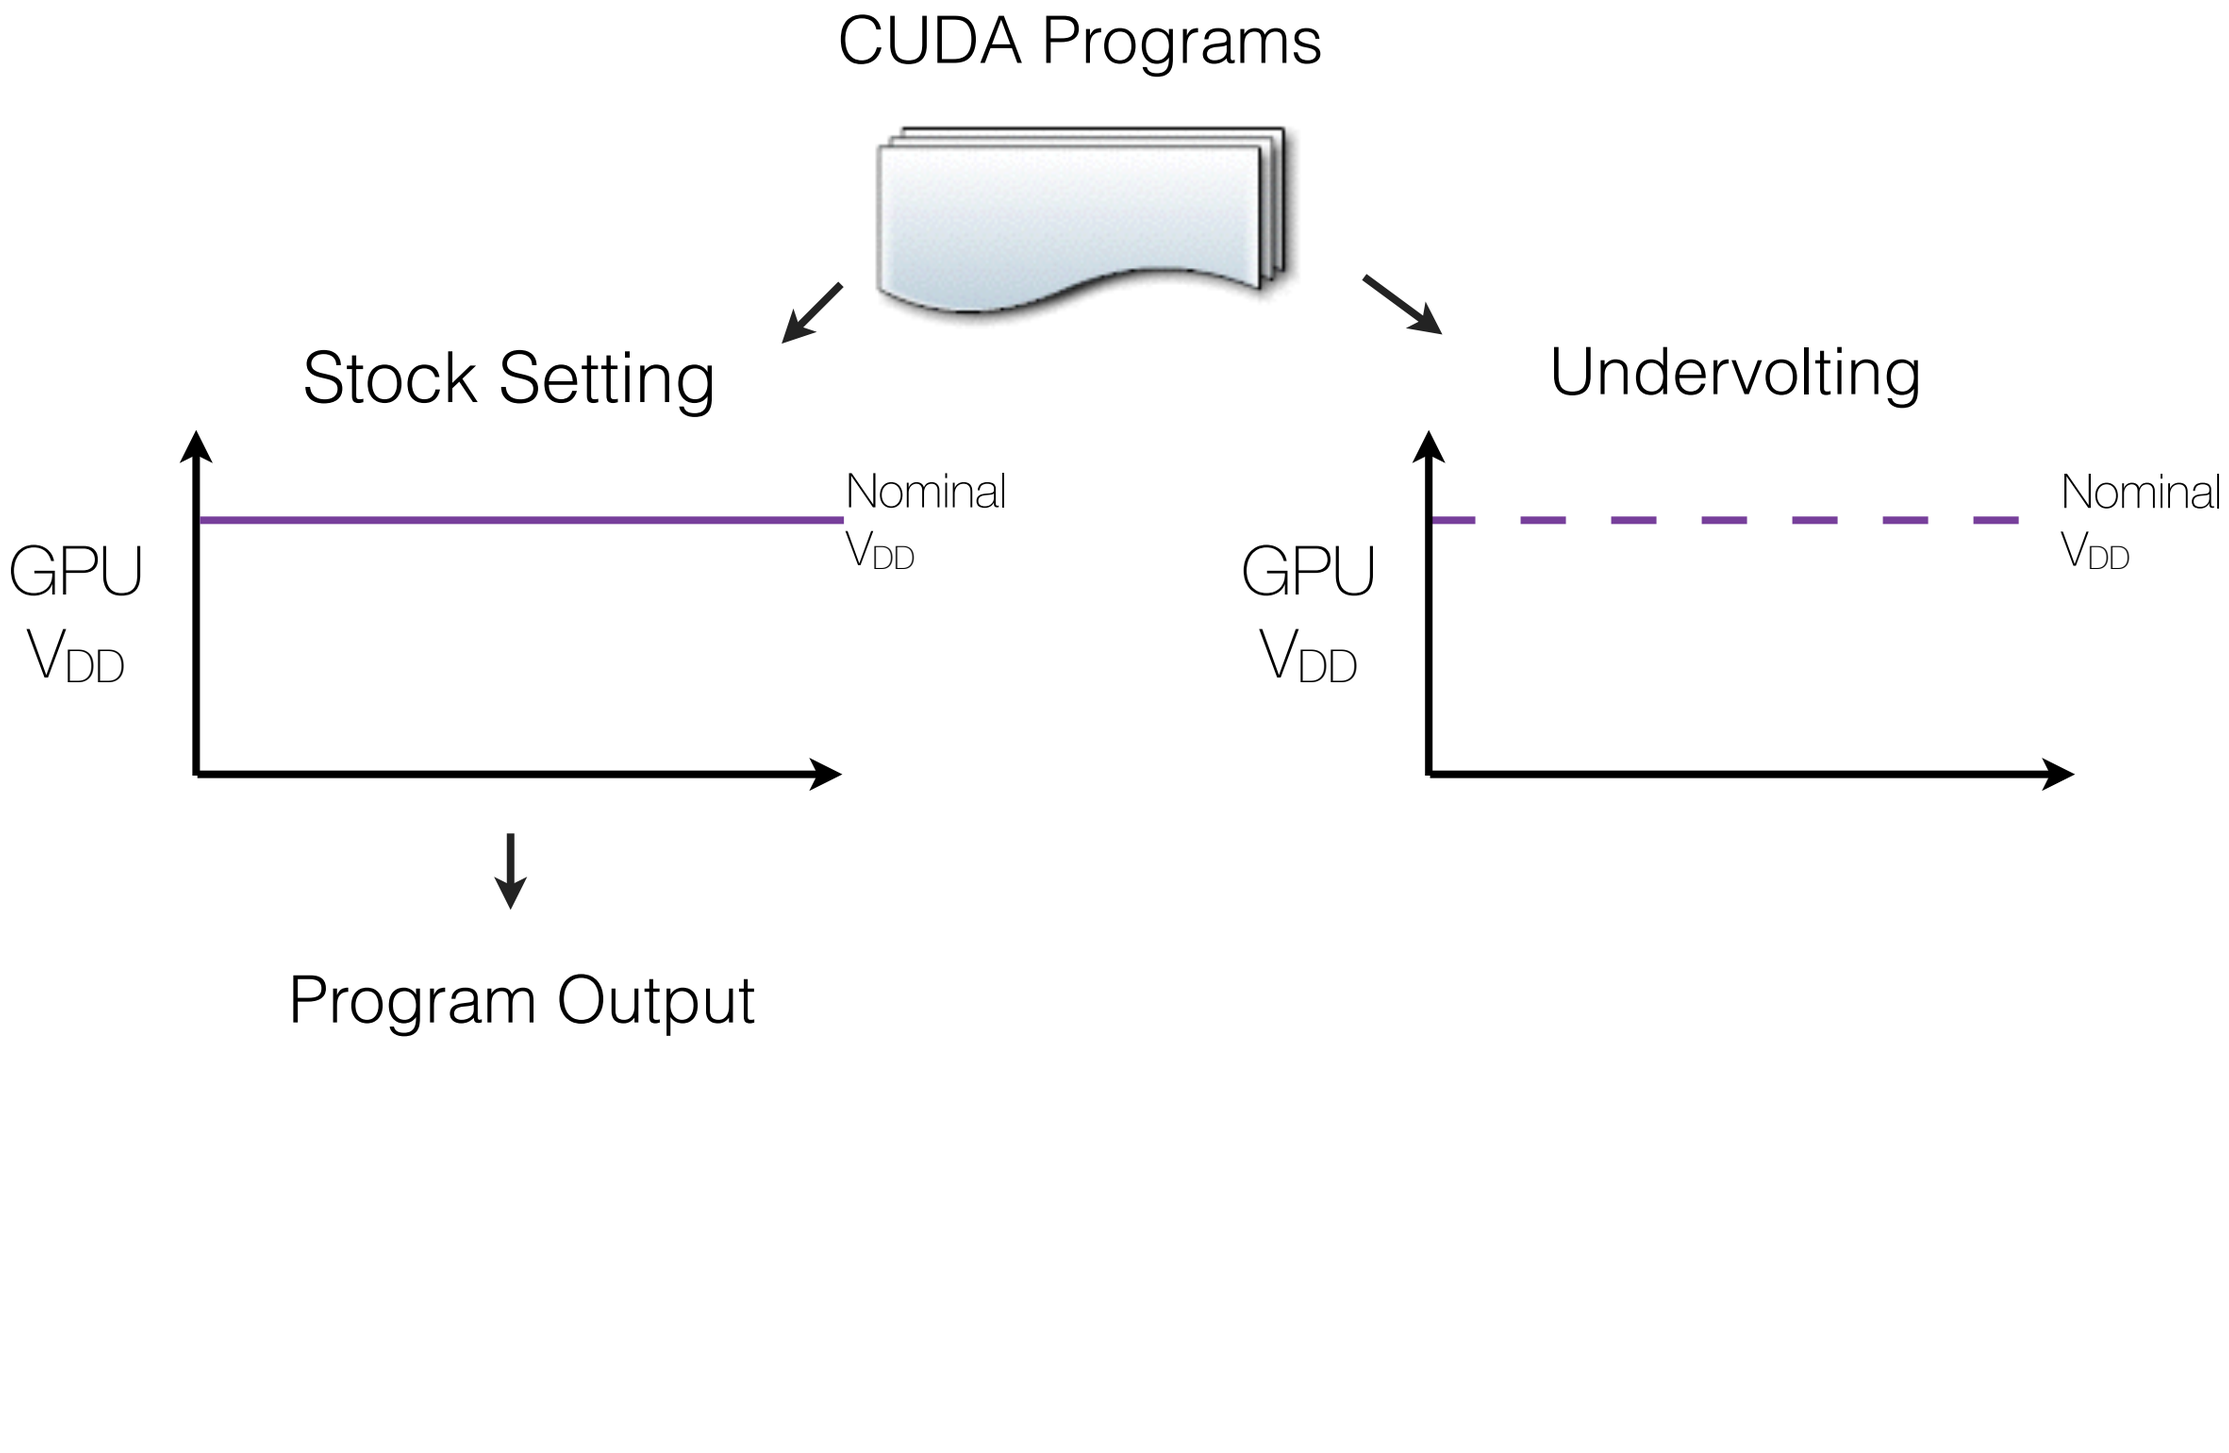

V_{min} Measurement

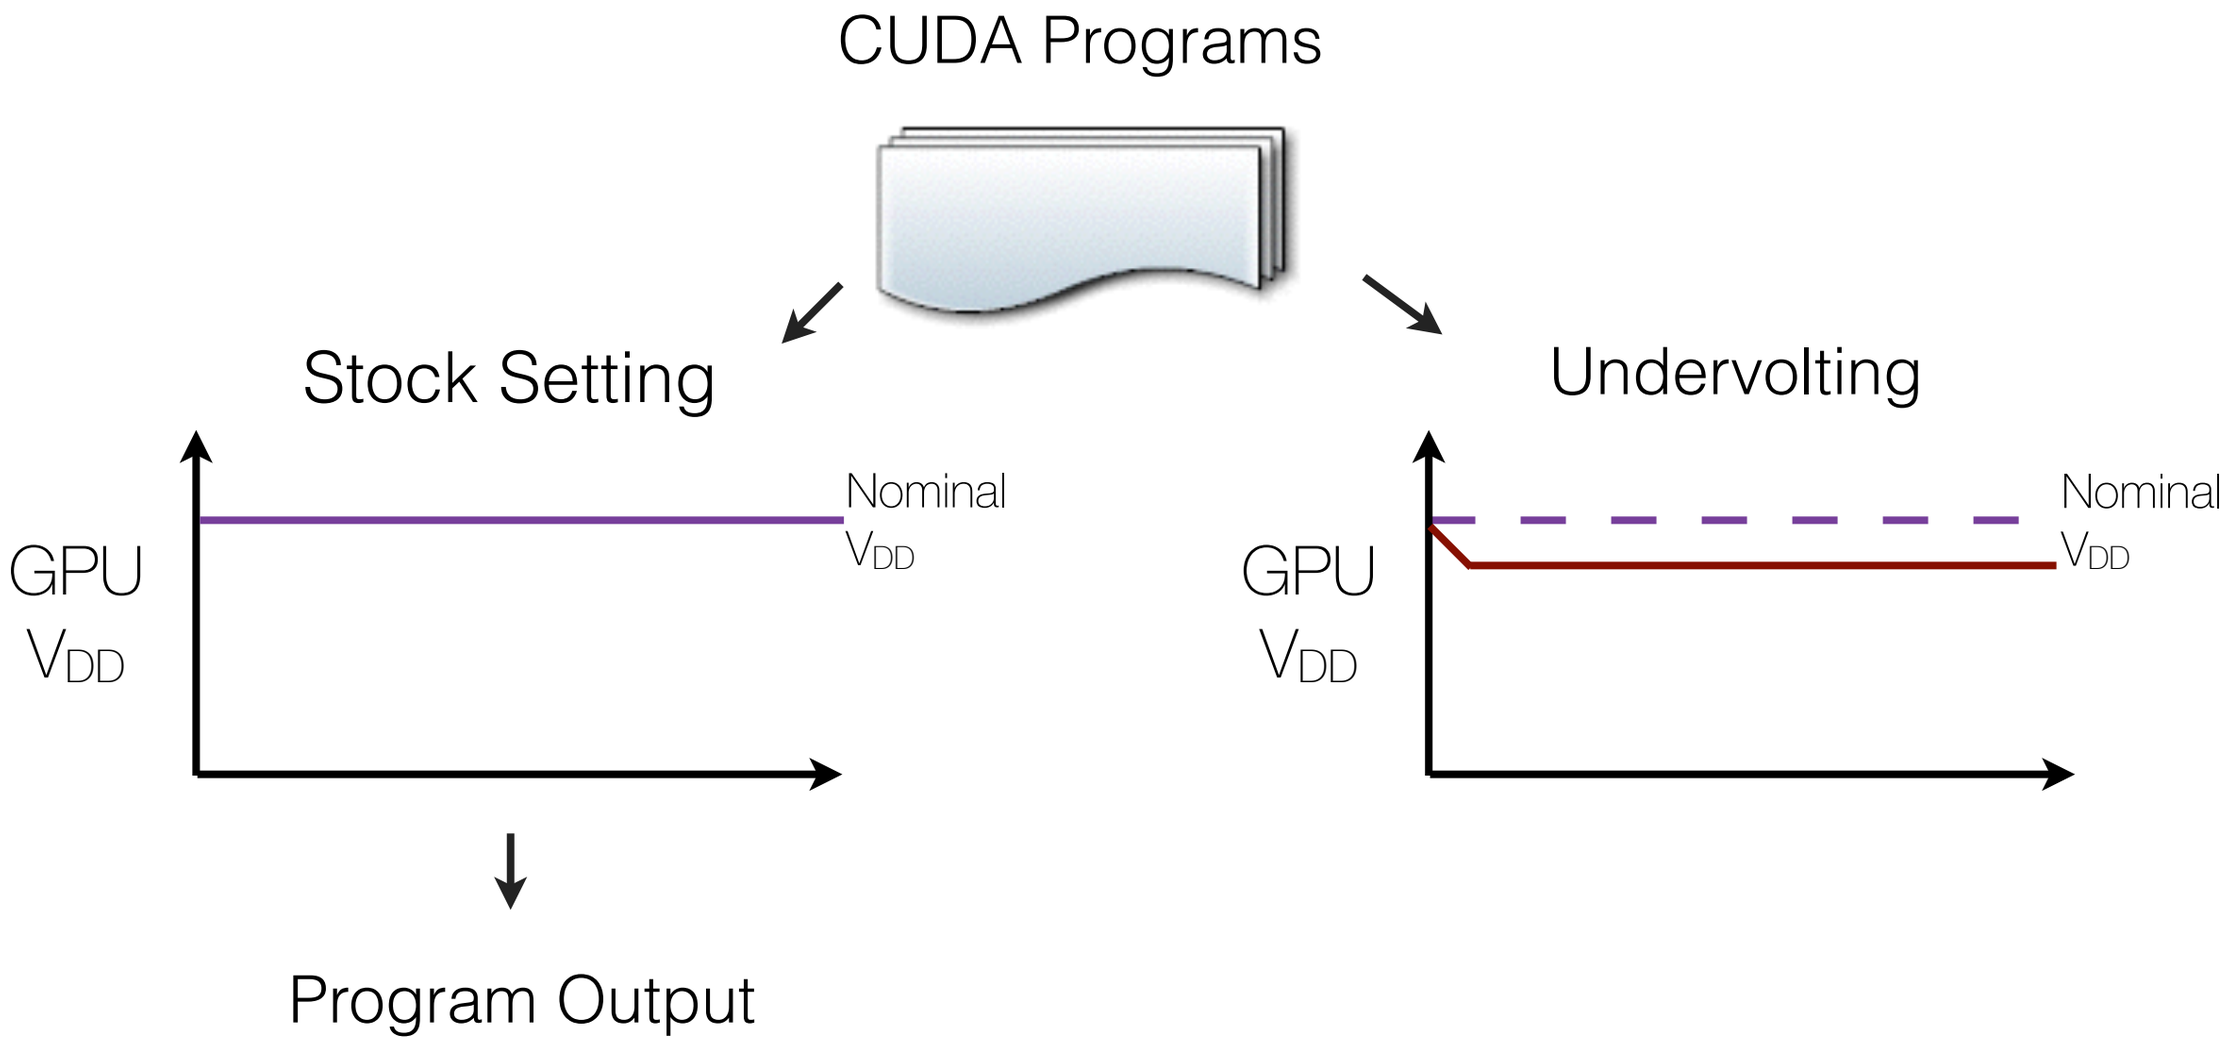

V_{min} Measurement

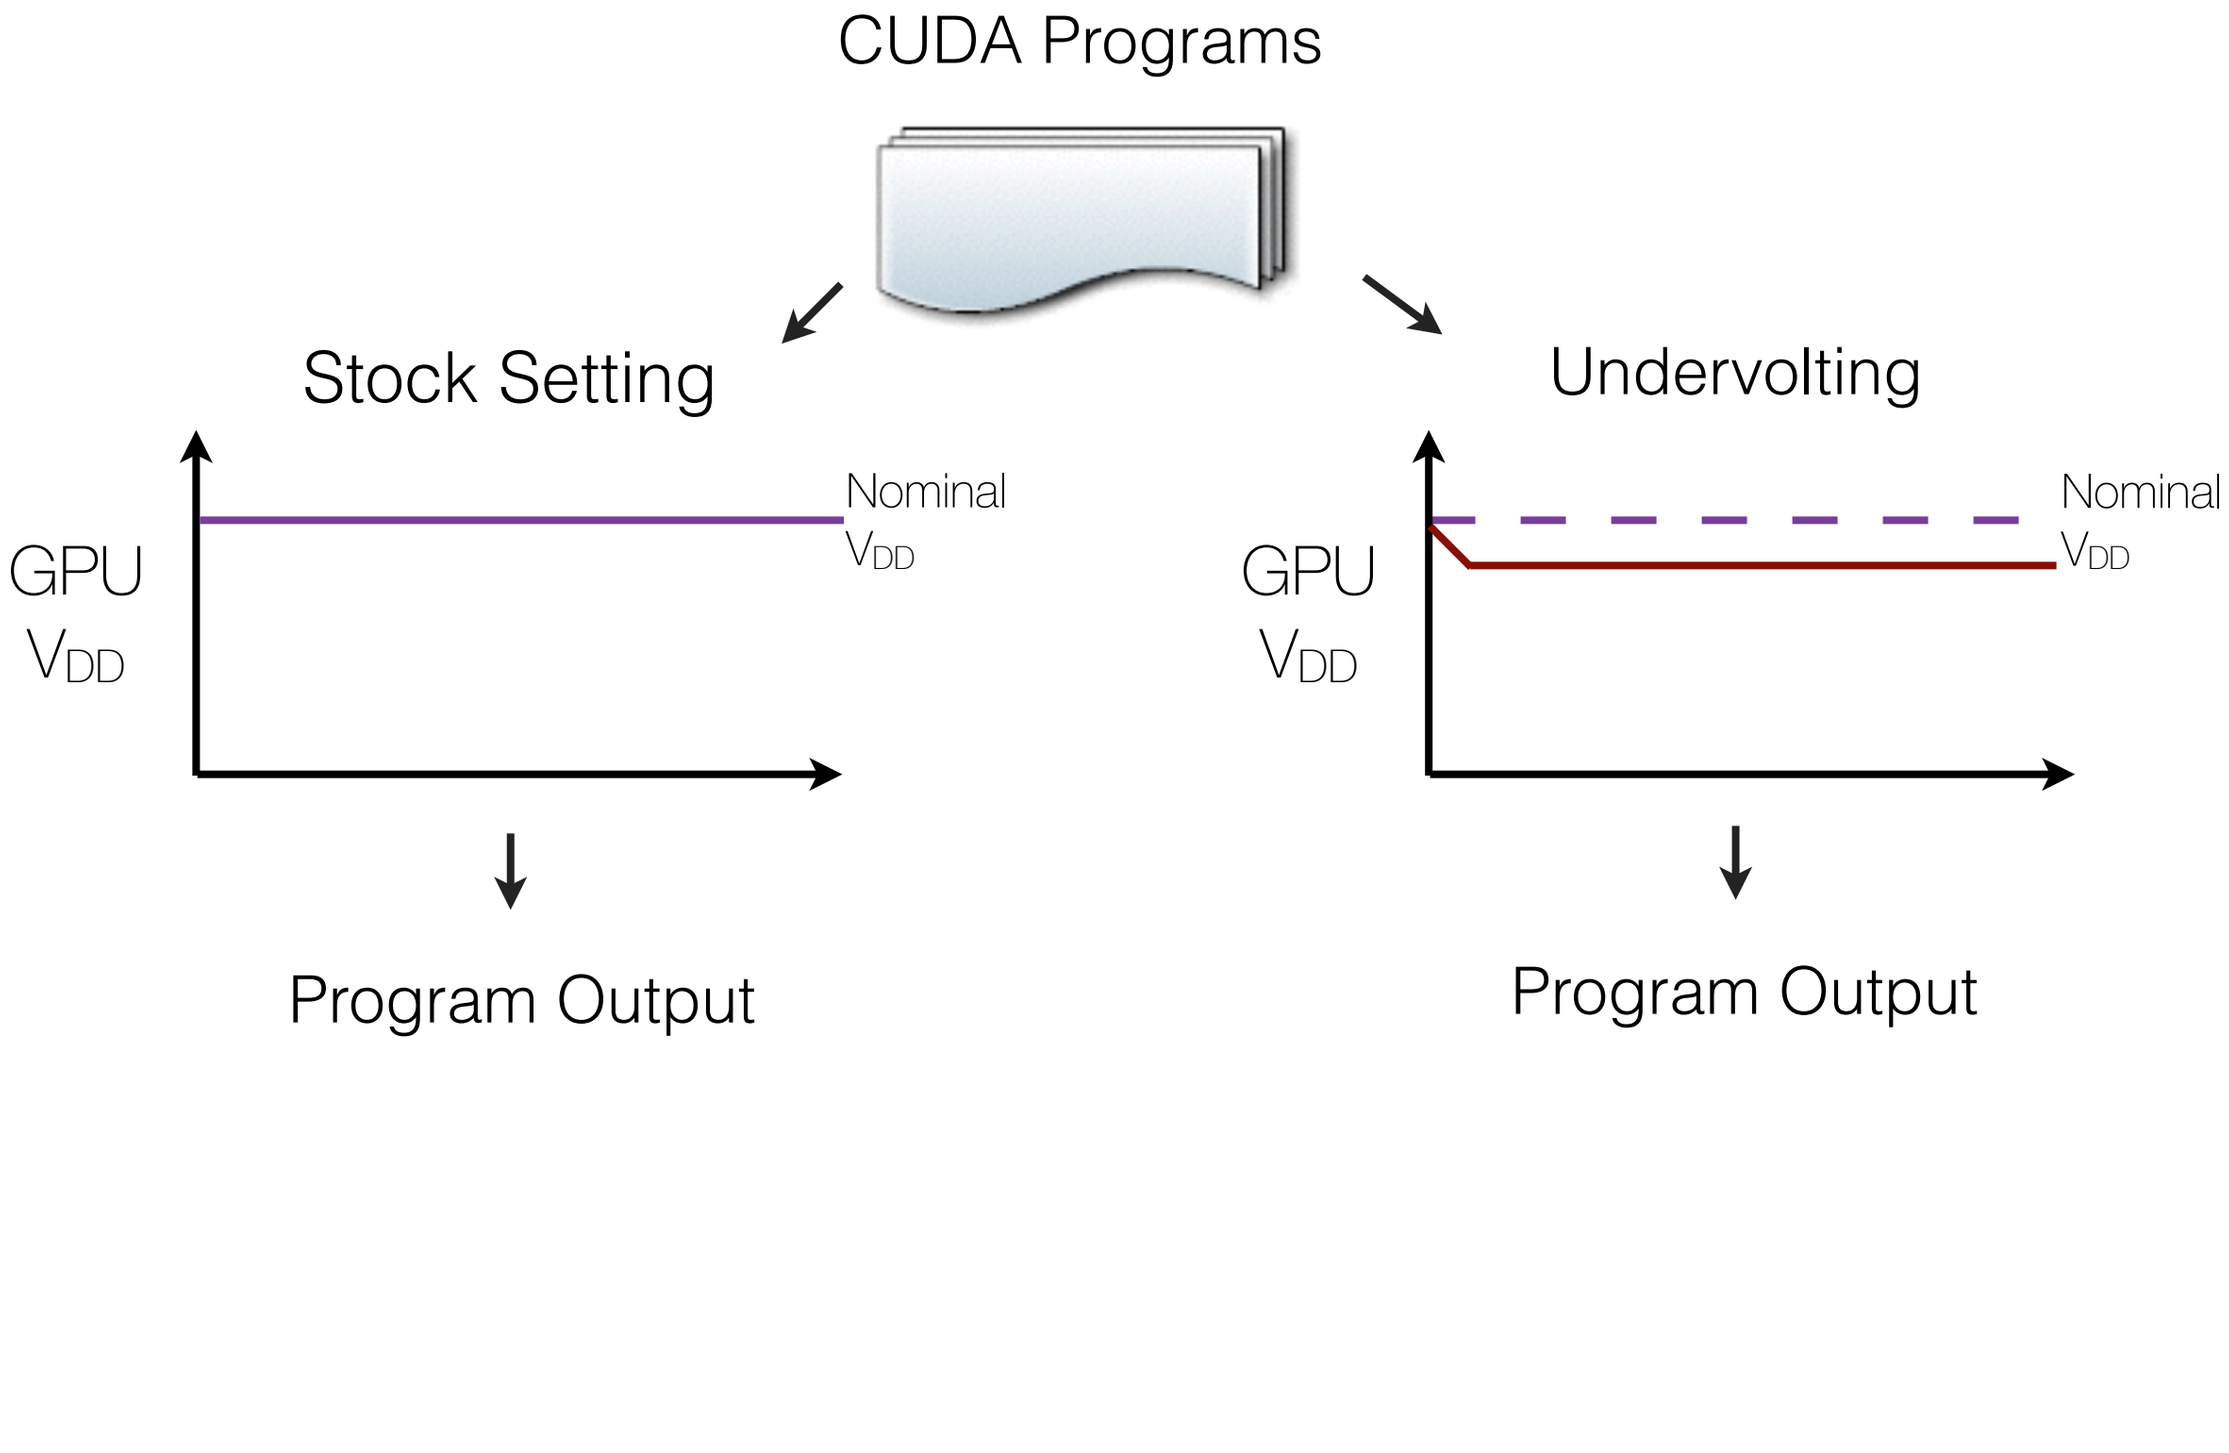

V_{min} Measurement

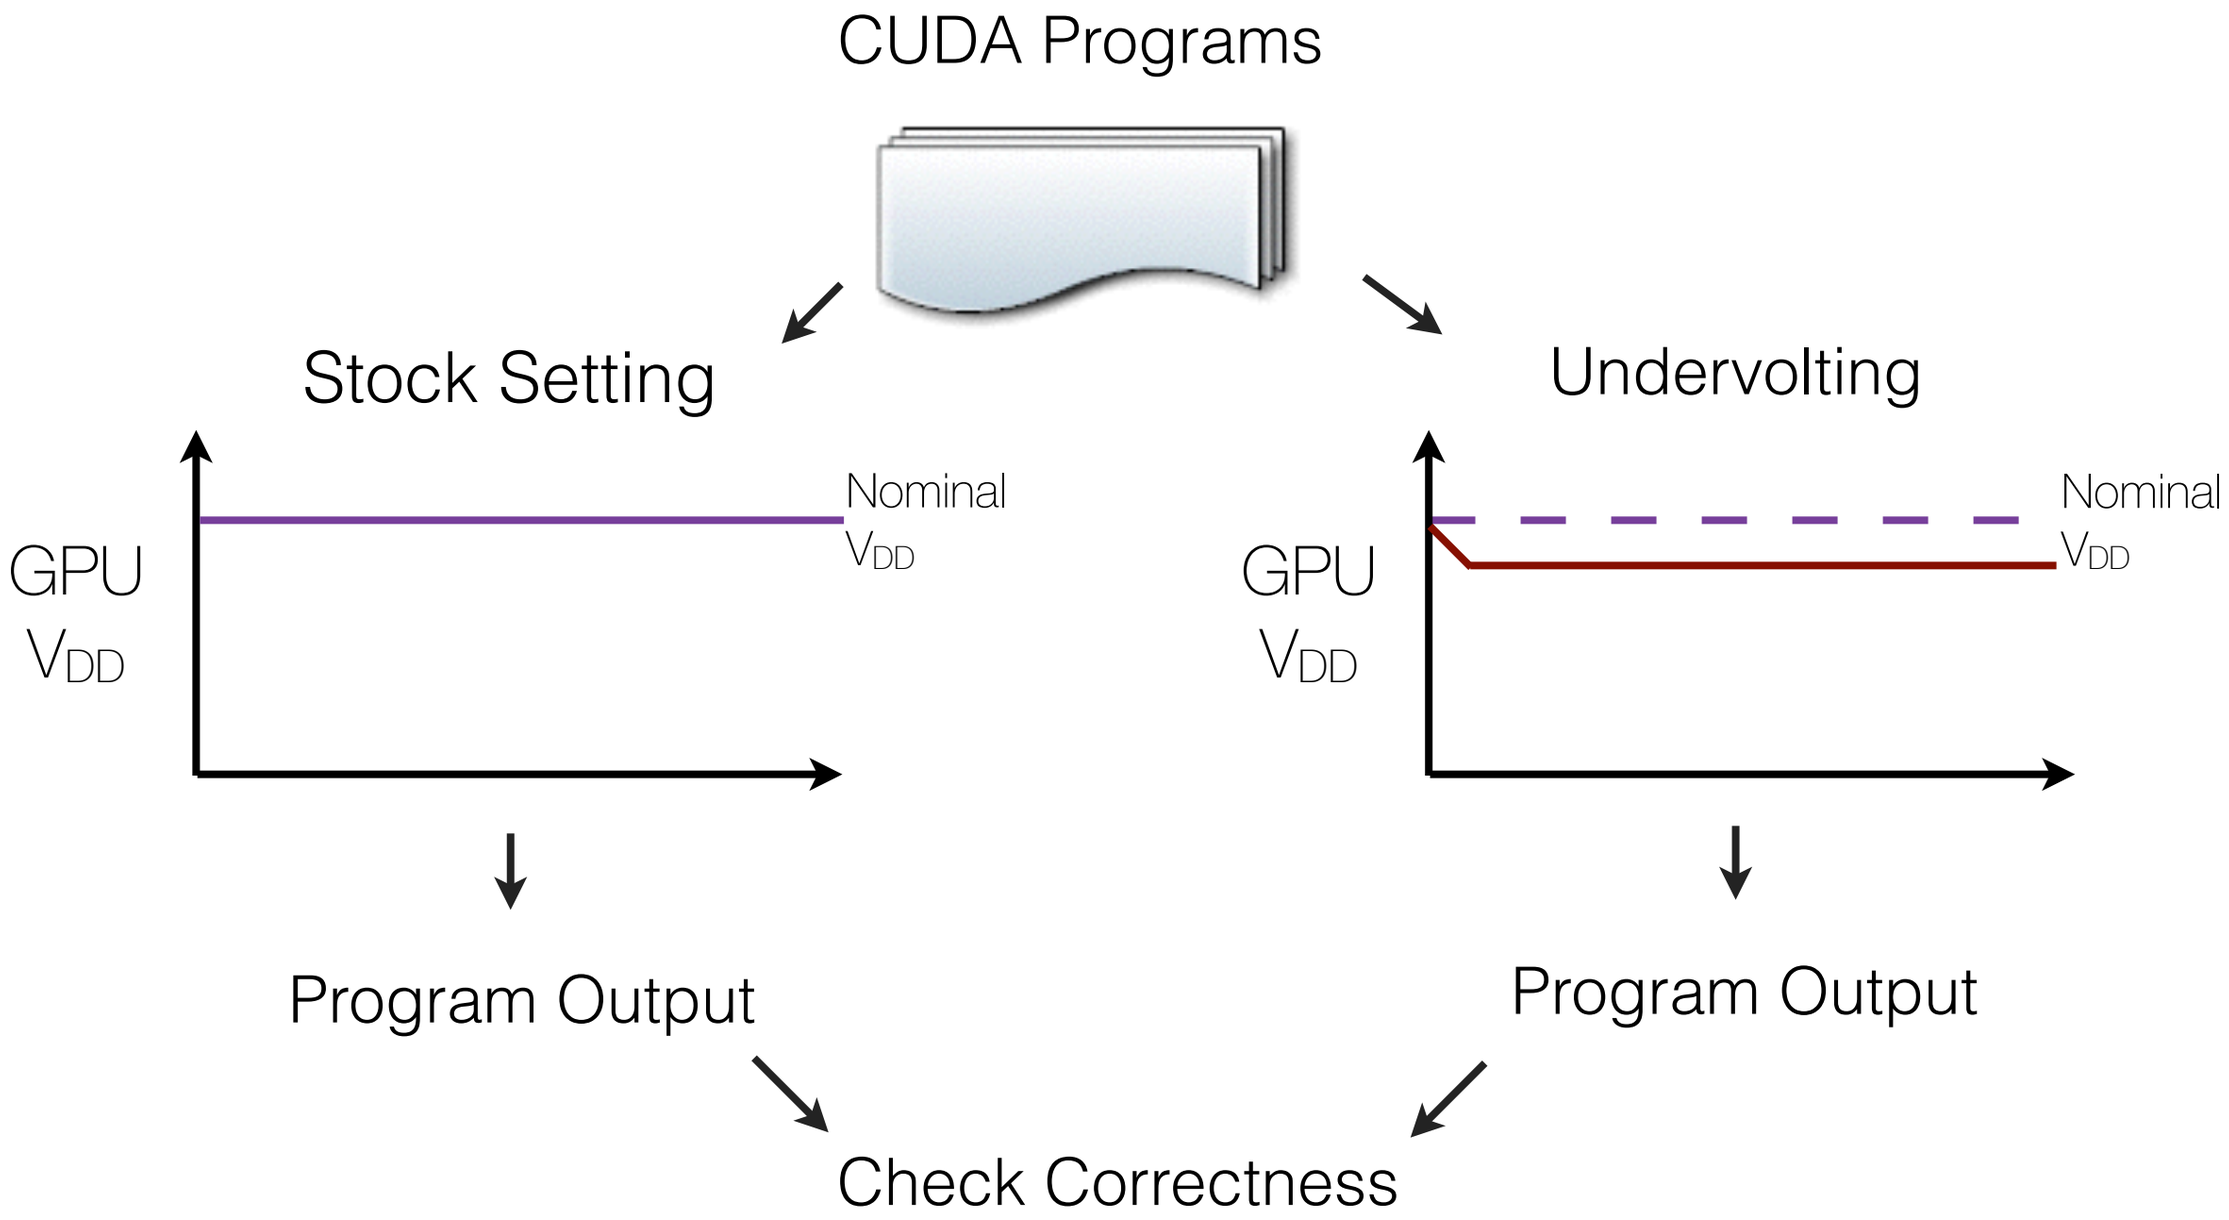

V_{min} Measurement

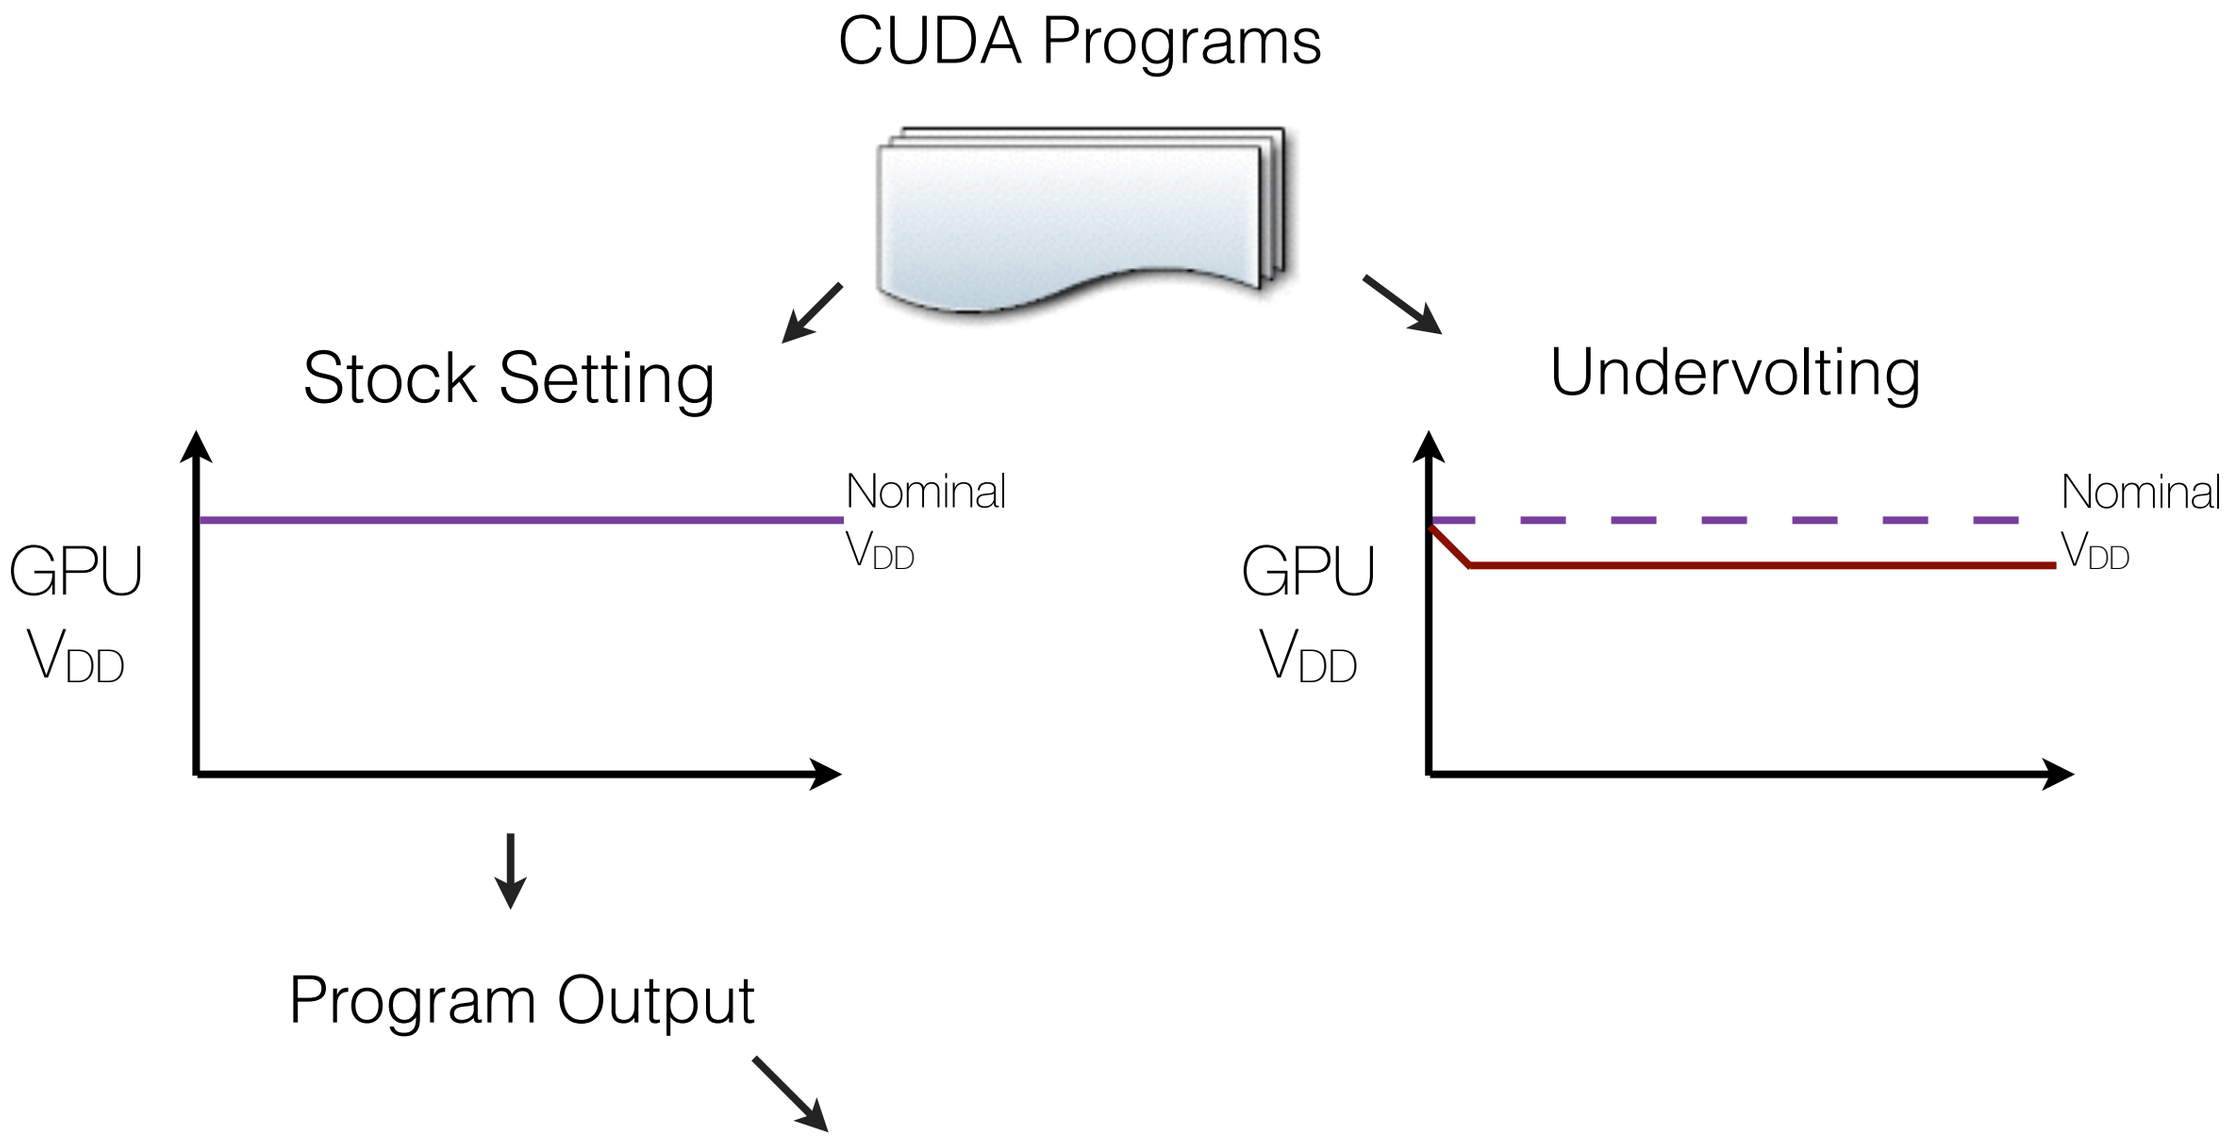

V_{min} Measurement

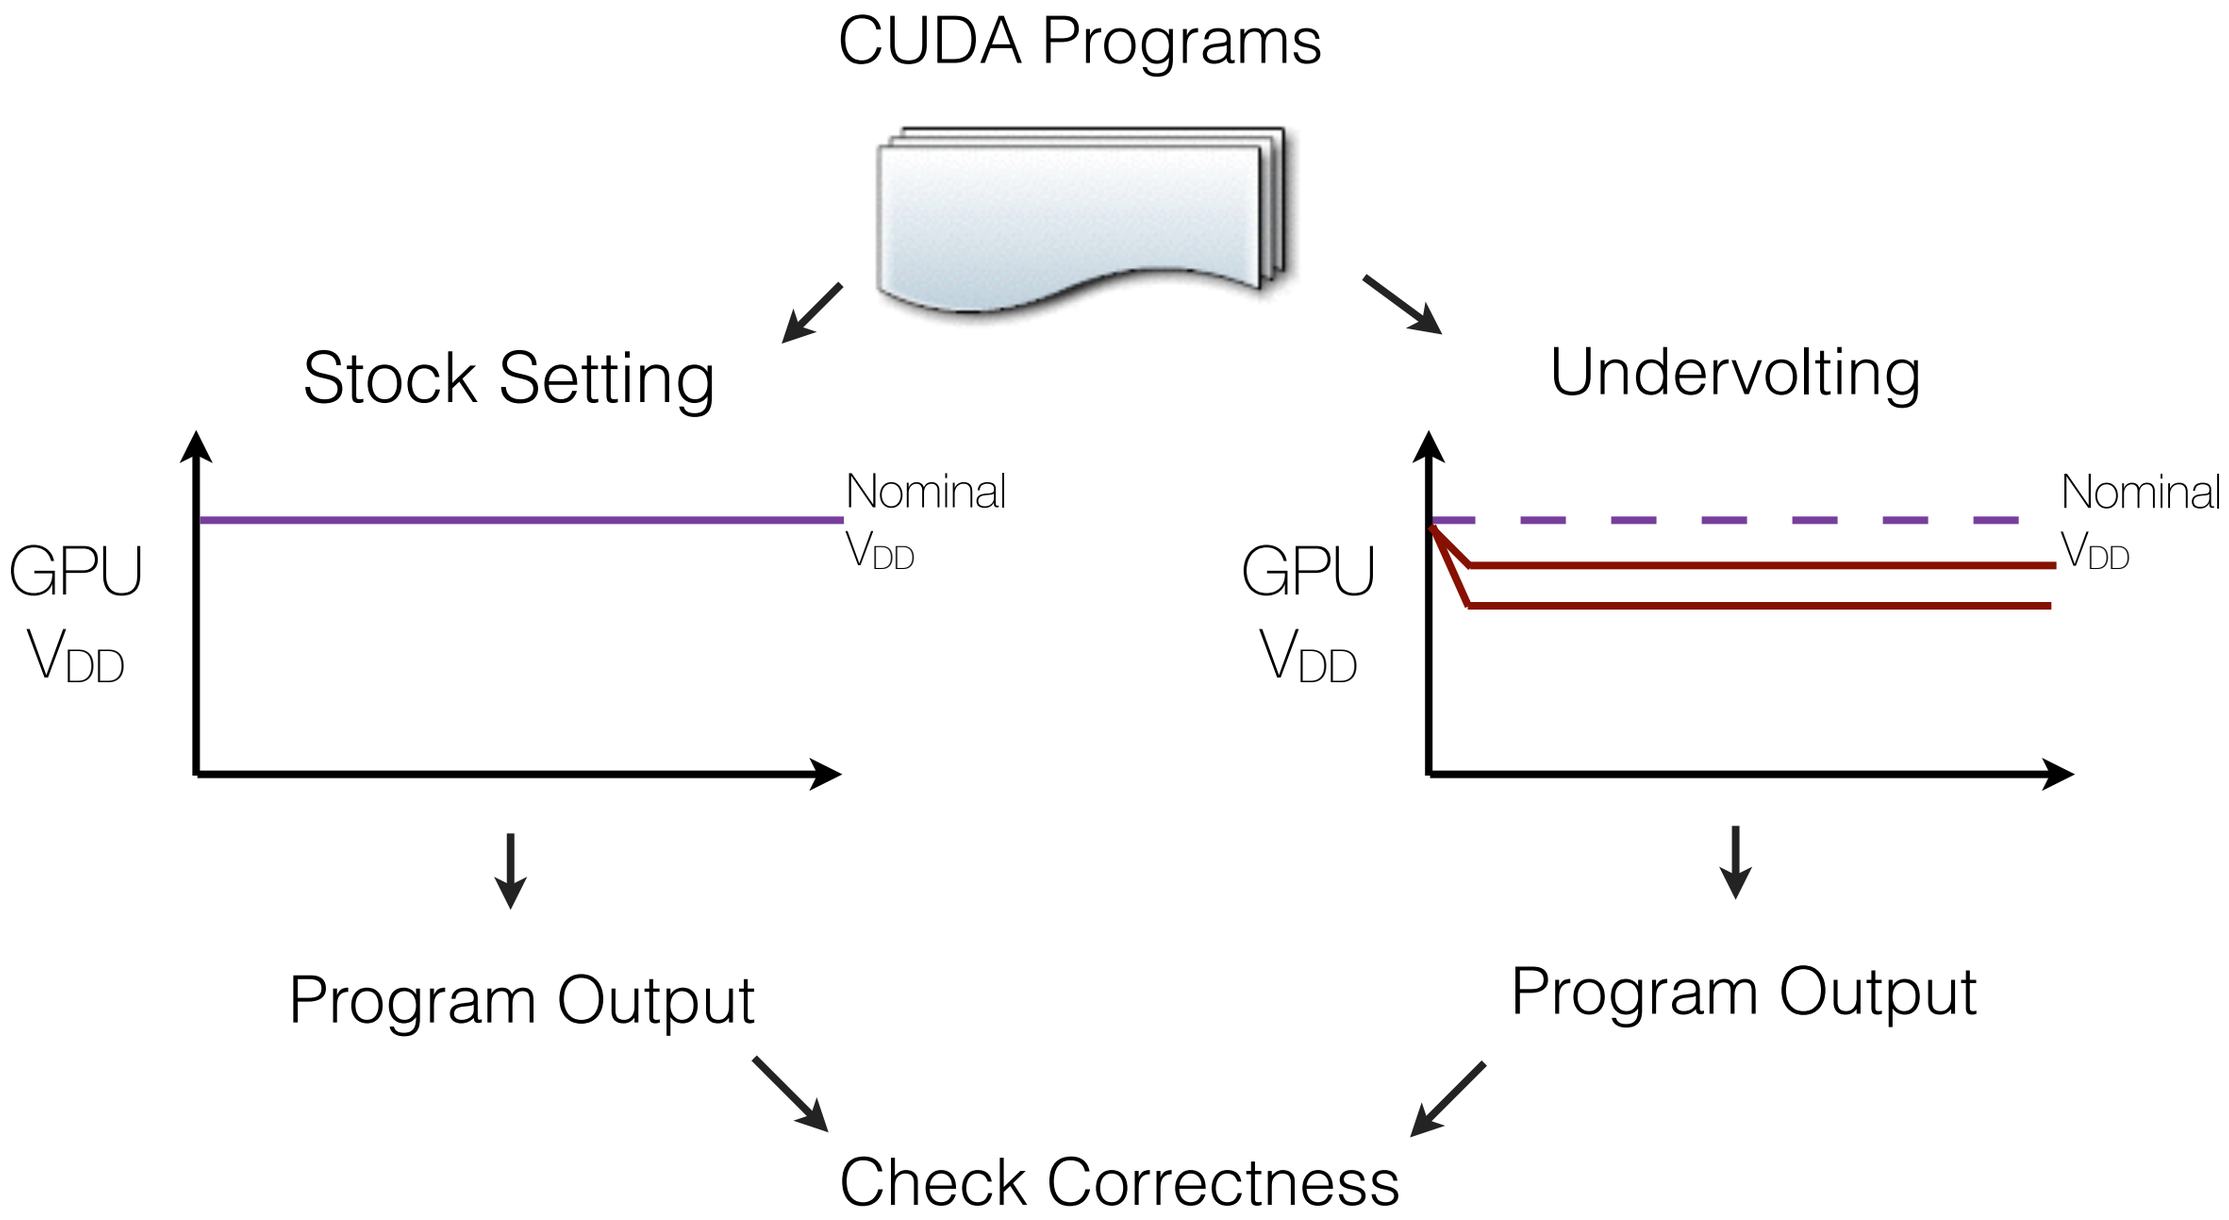

V_{min} Measurement

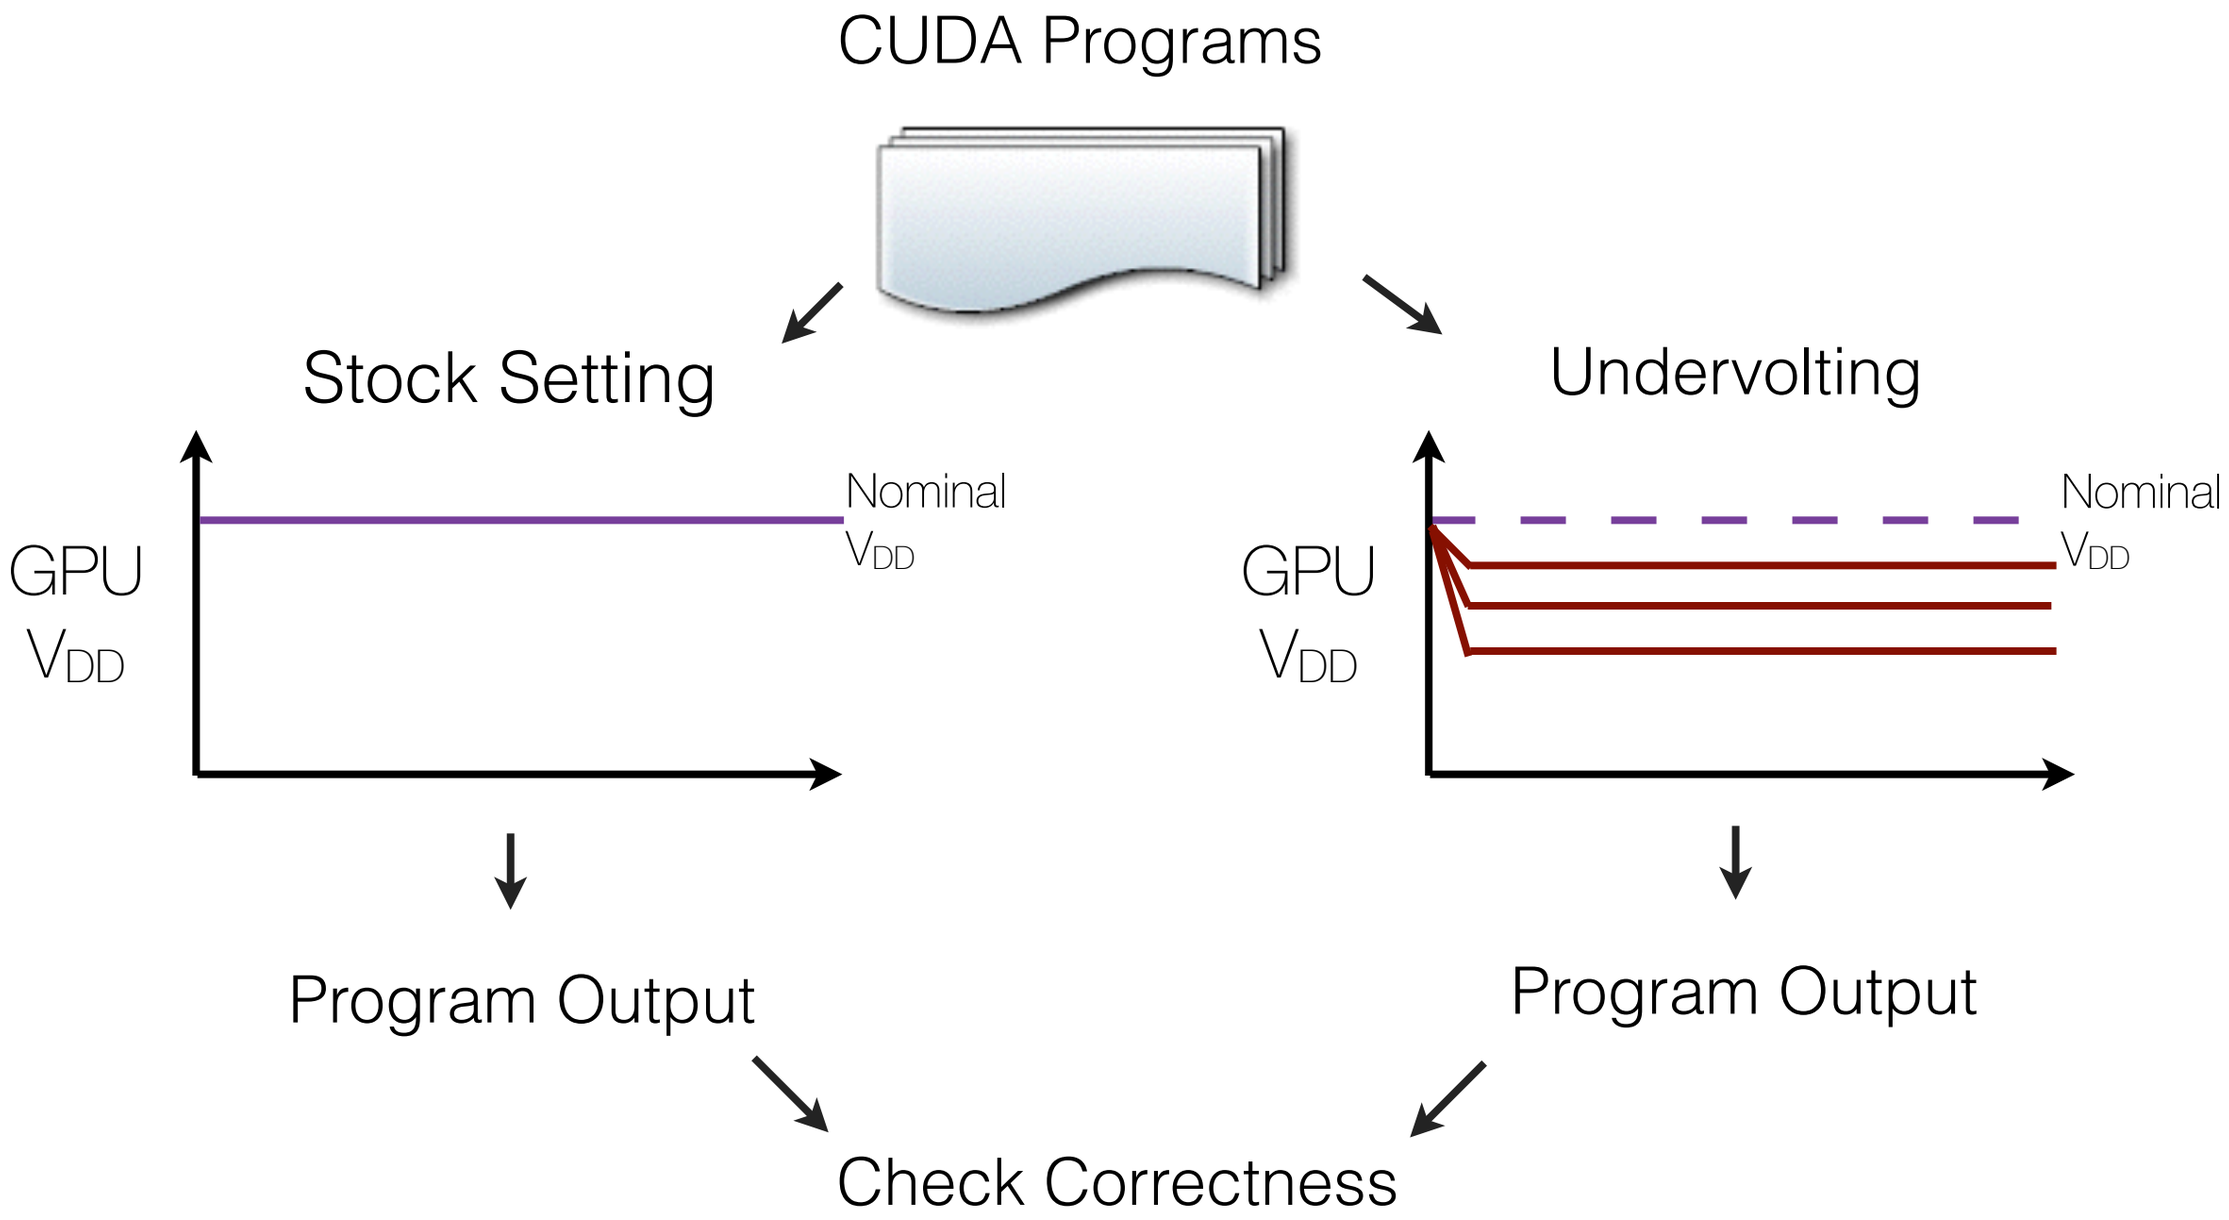

V_{min} Measurement

CUDA Programs

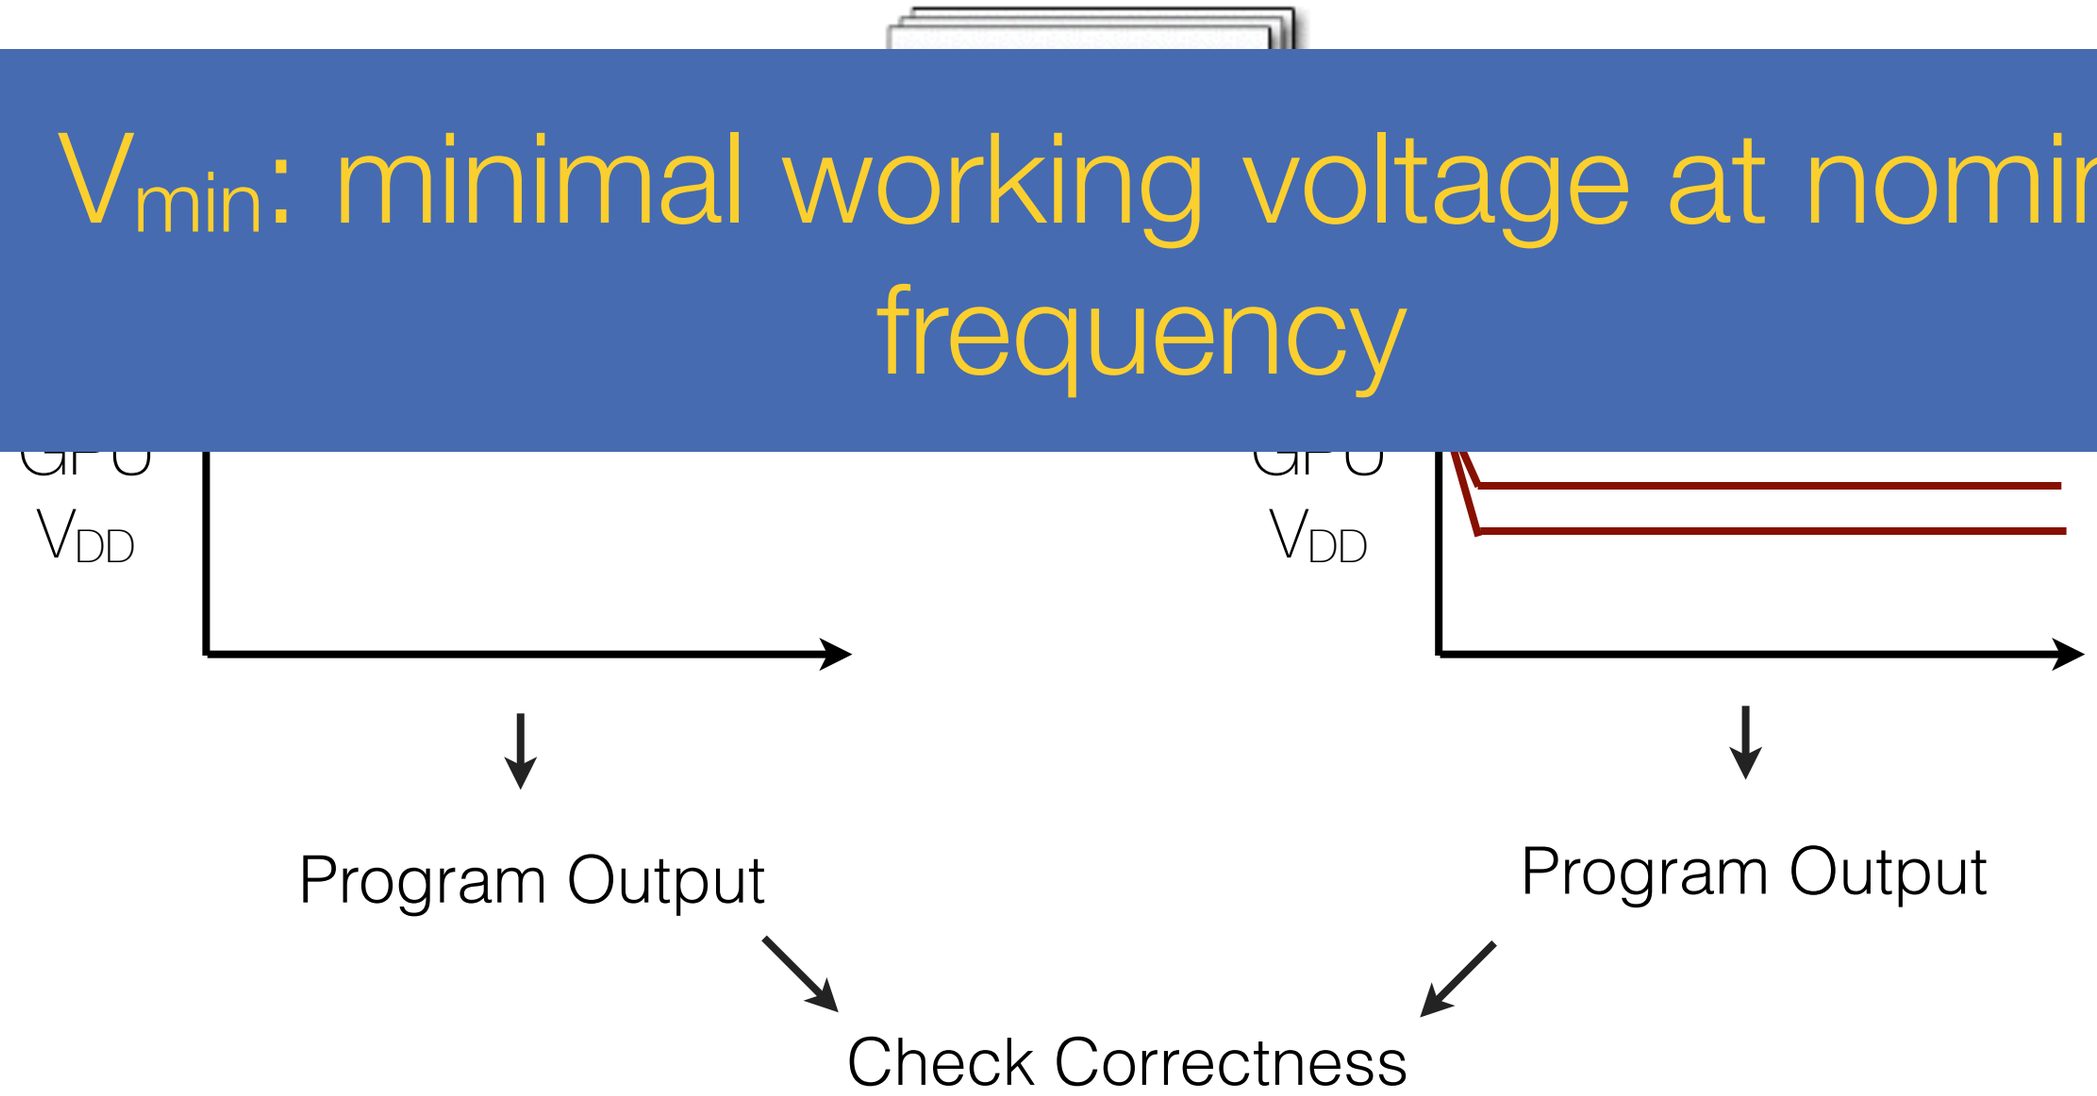

V_{min}: minimal working voltage at nominal frequency

Measurement Results

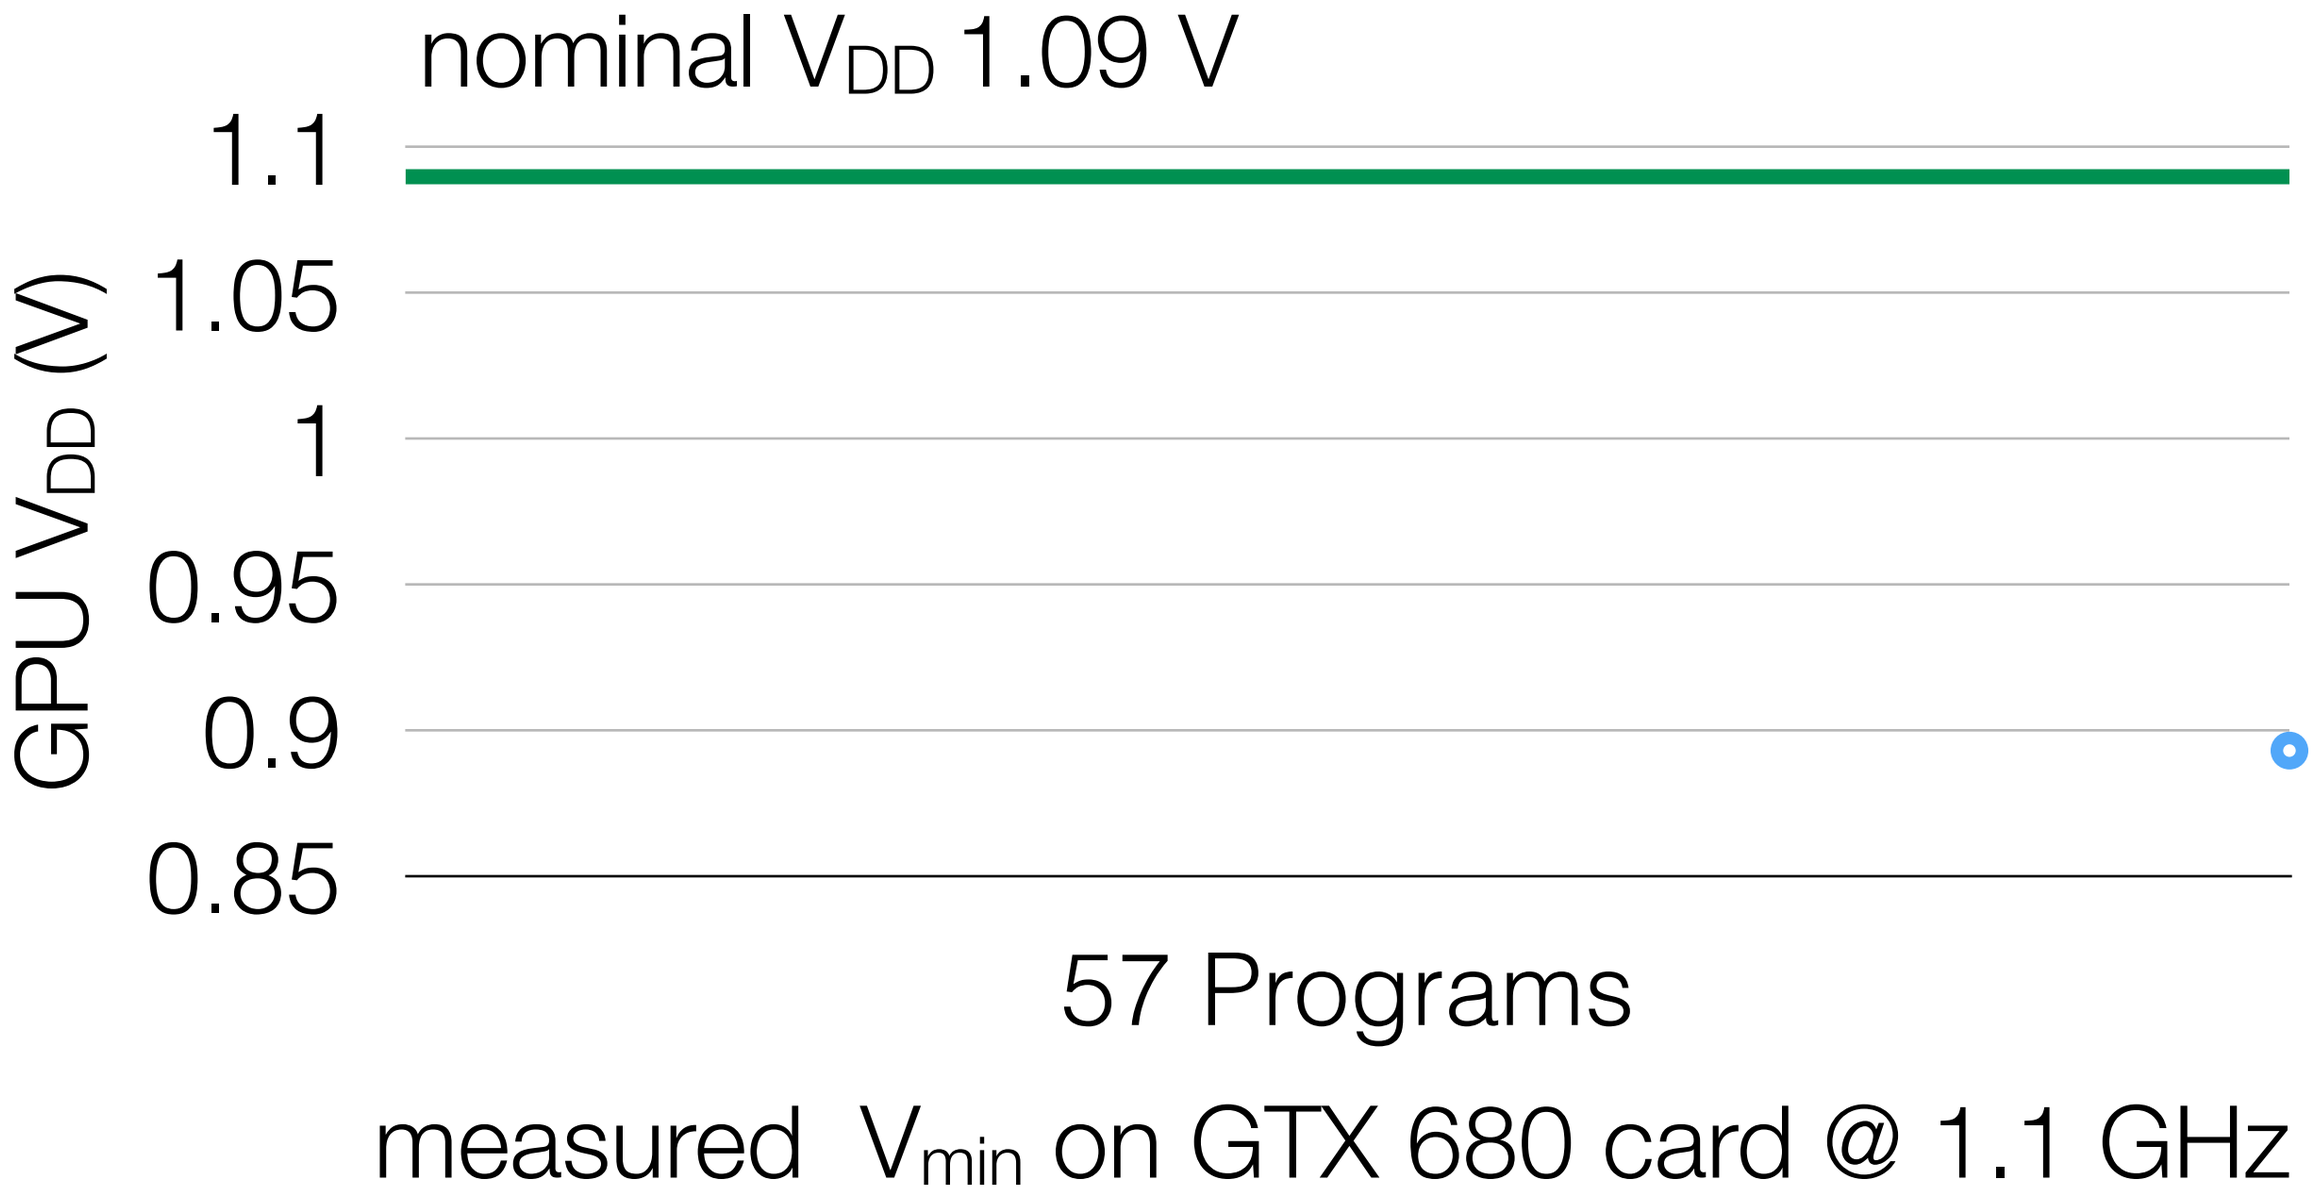

Measurement Results

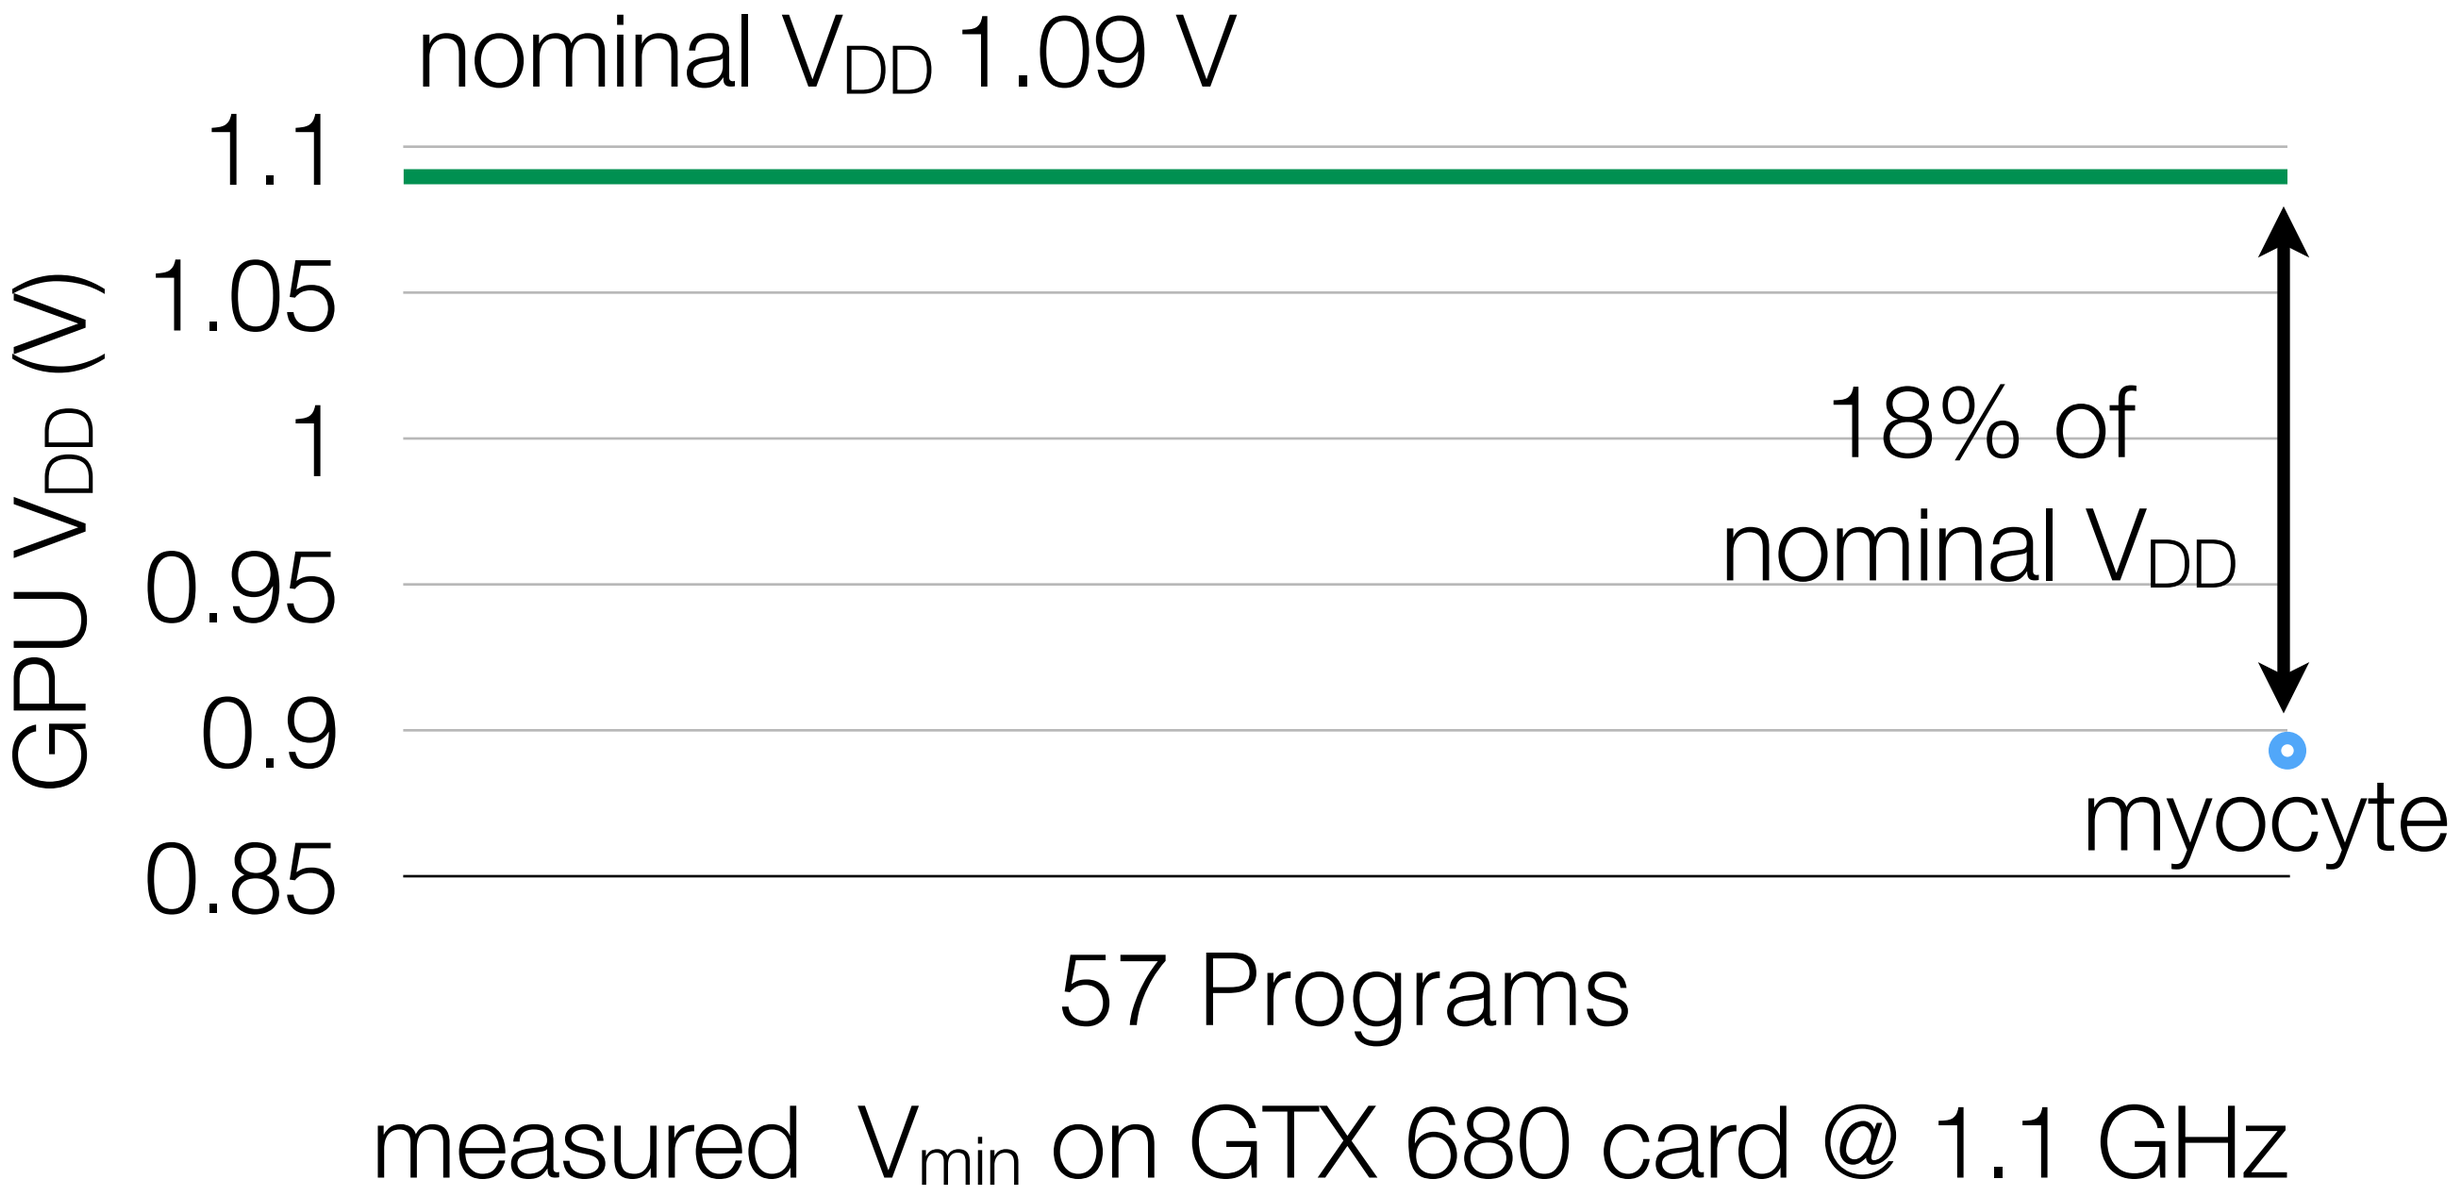

Measurement Results

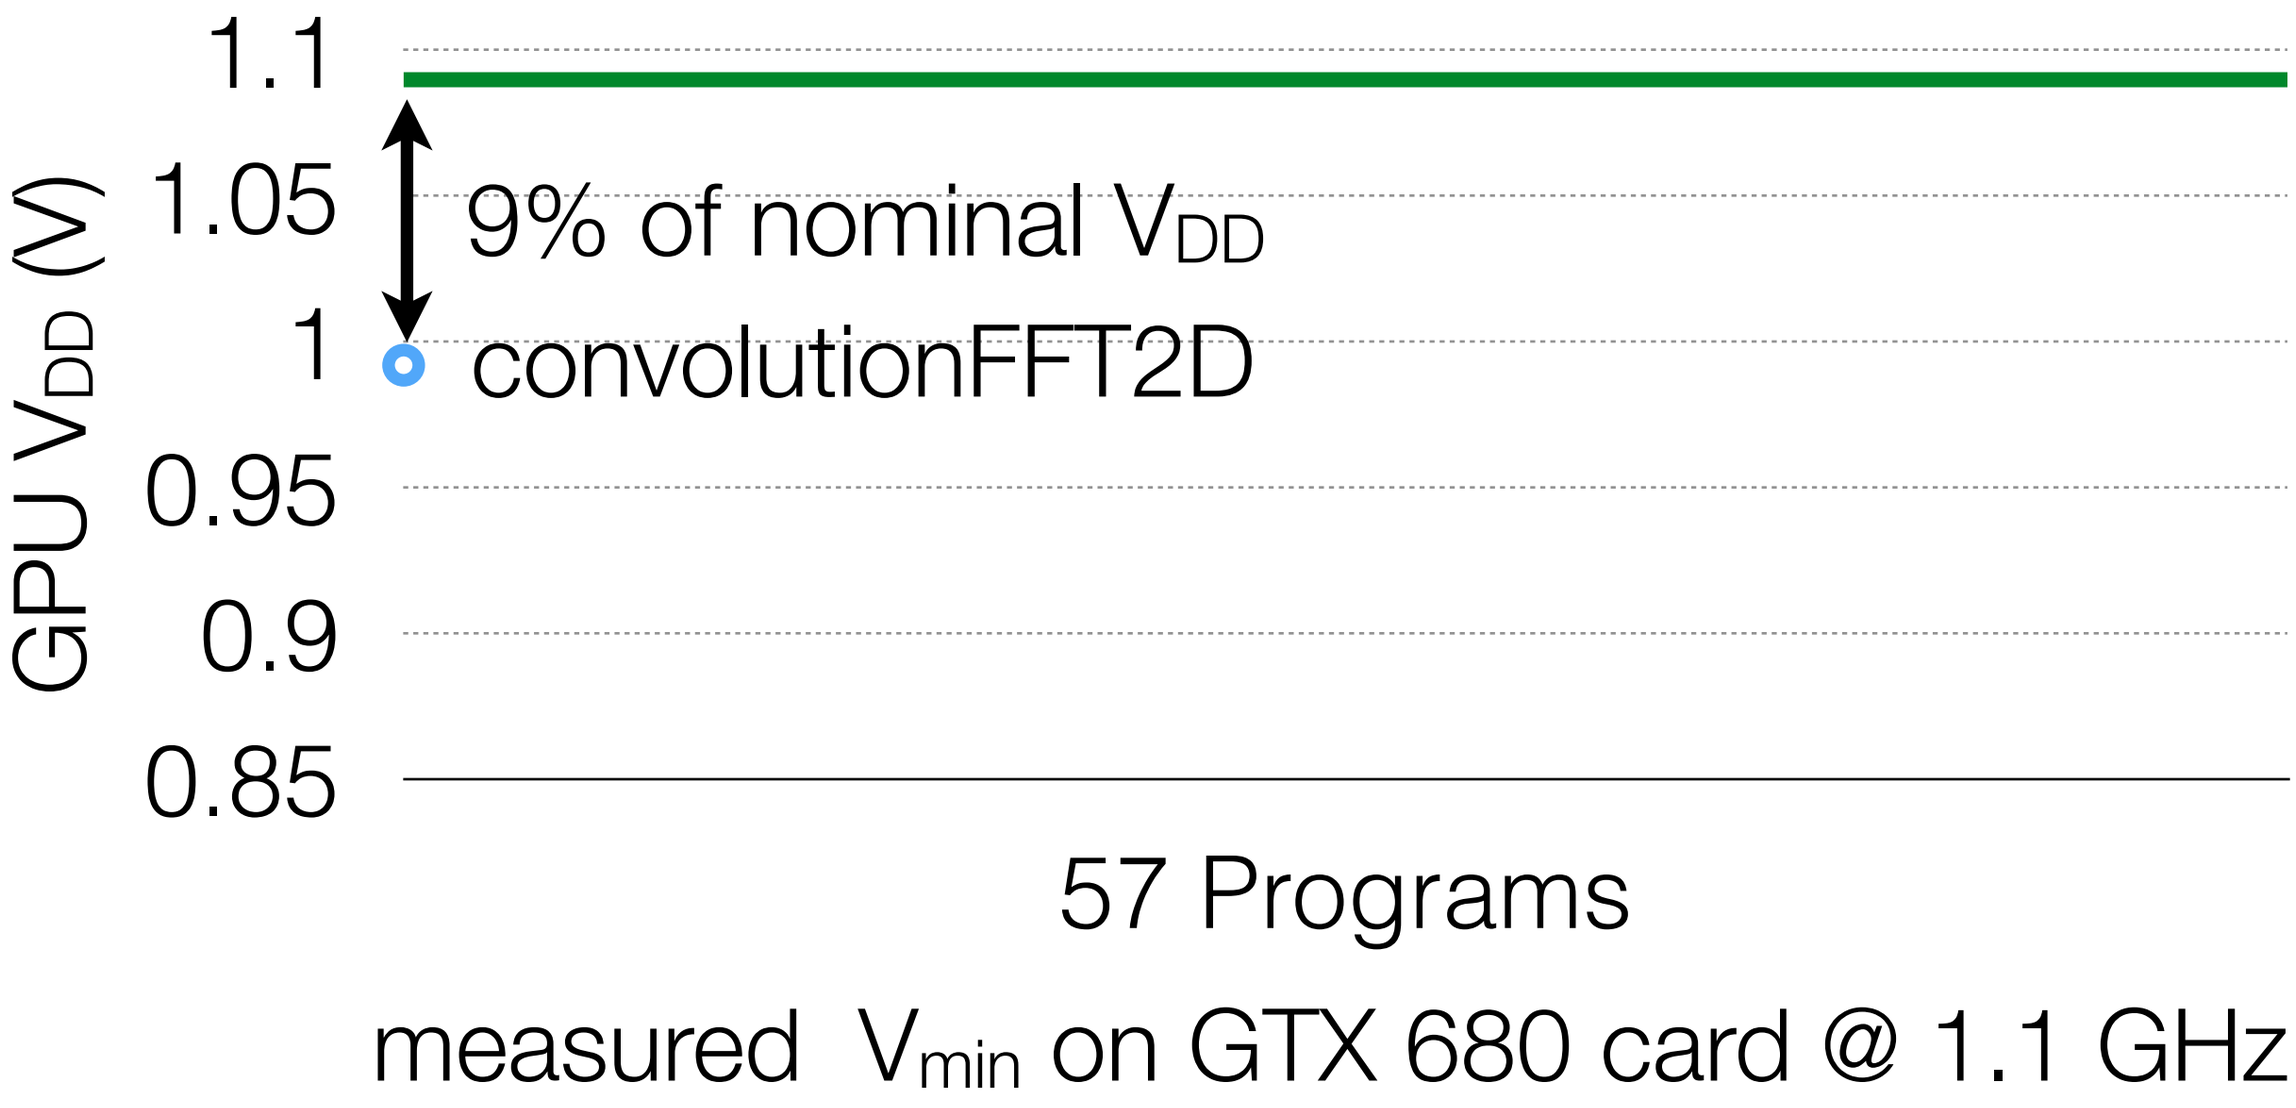

Measurement Results

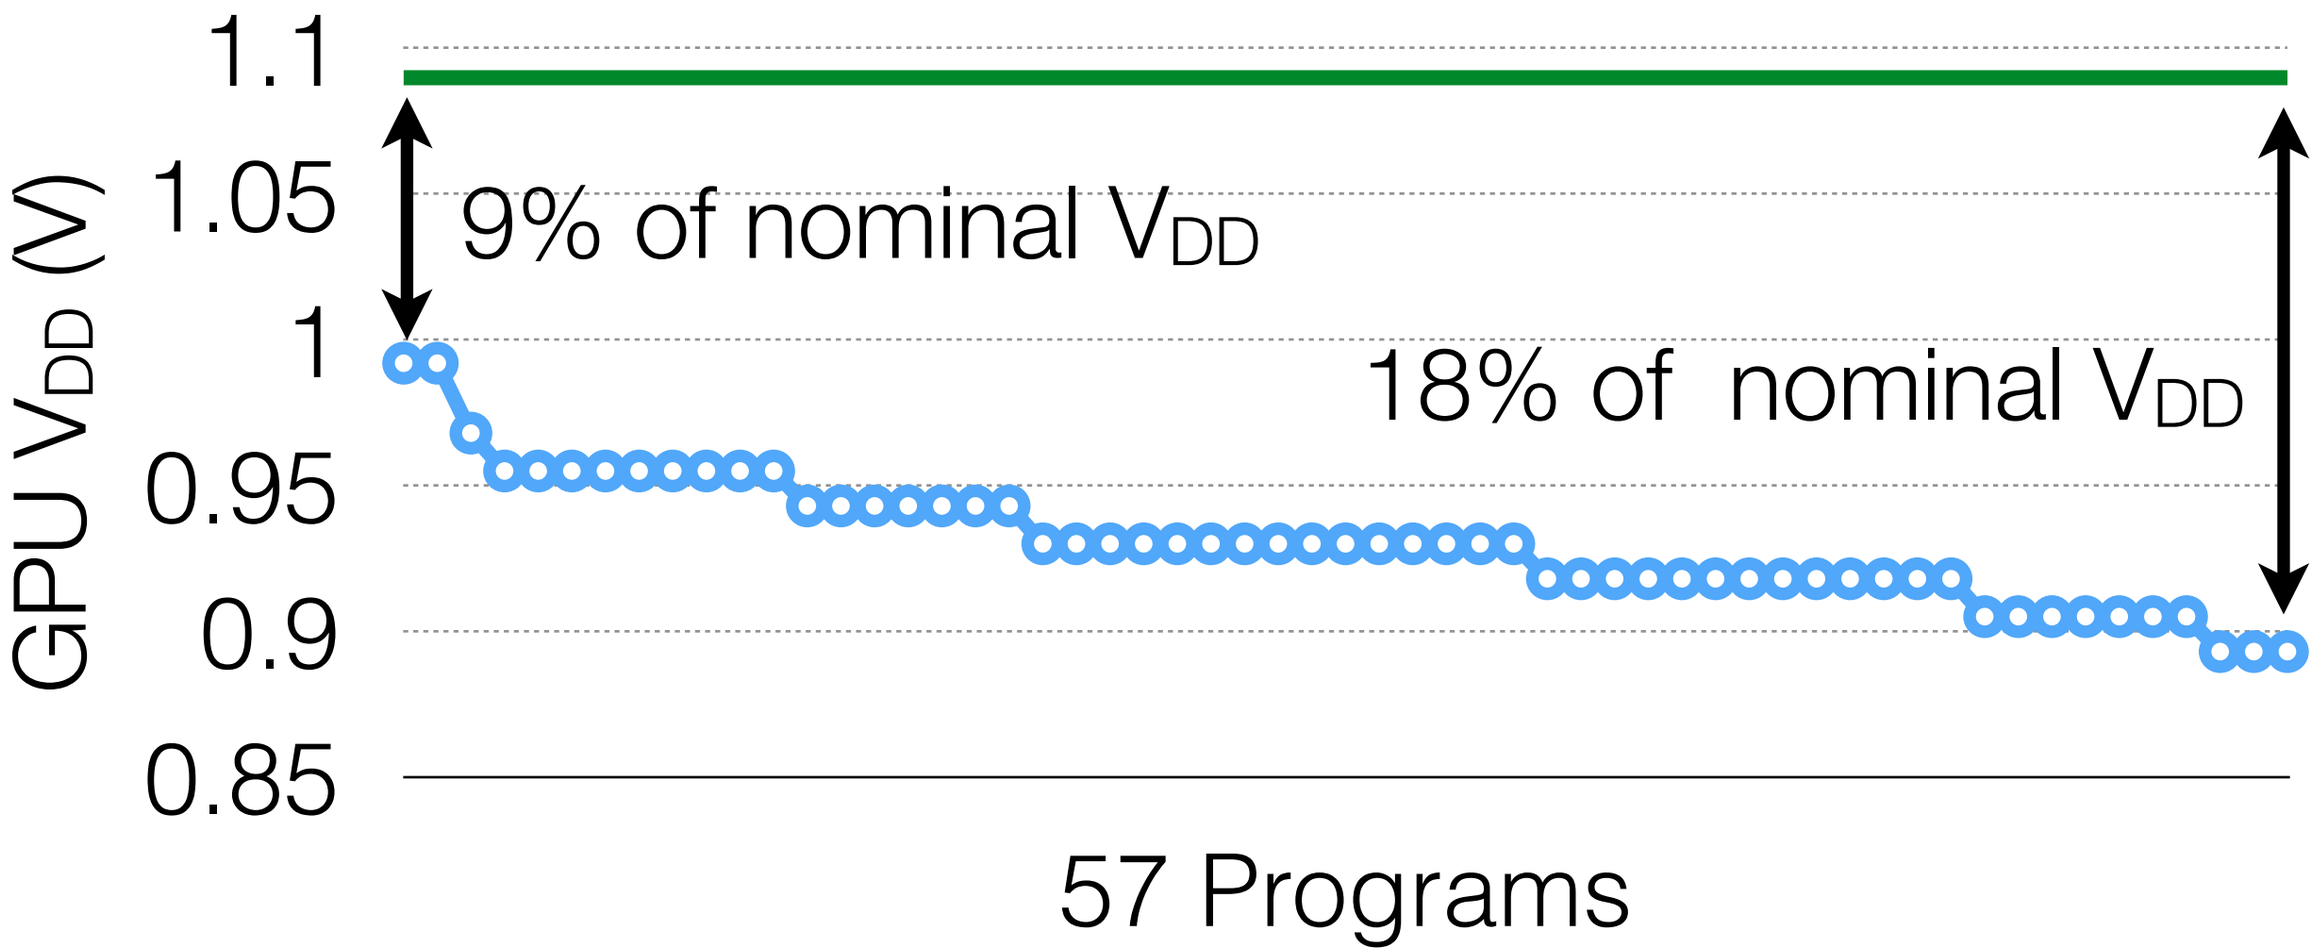

Measurement Results

- Voltage guardband: 9% - 18%

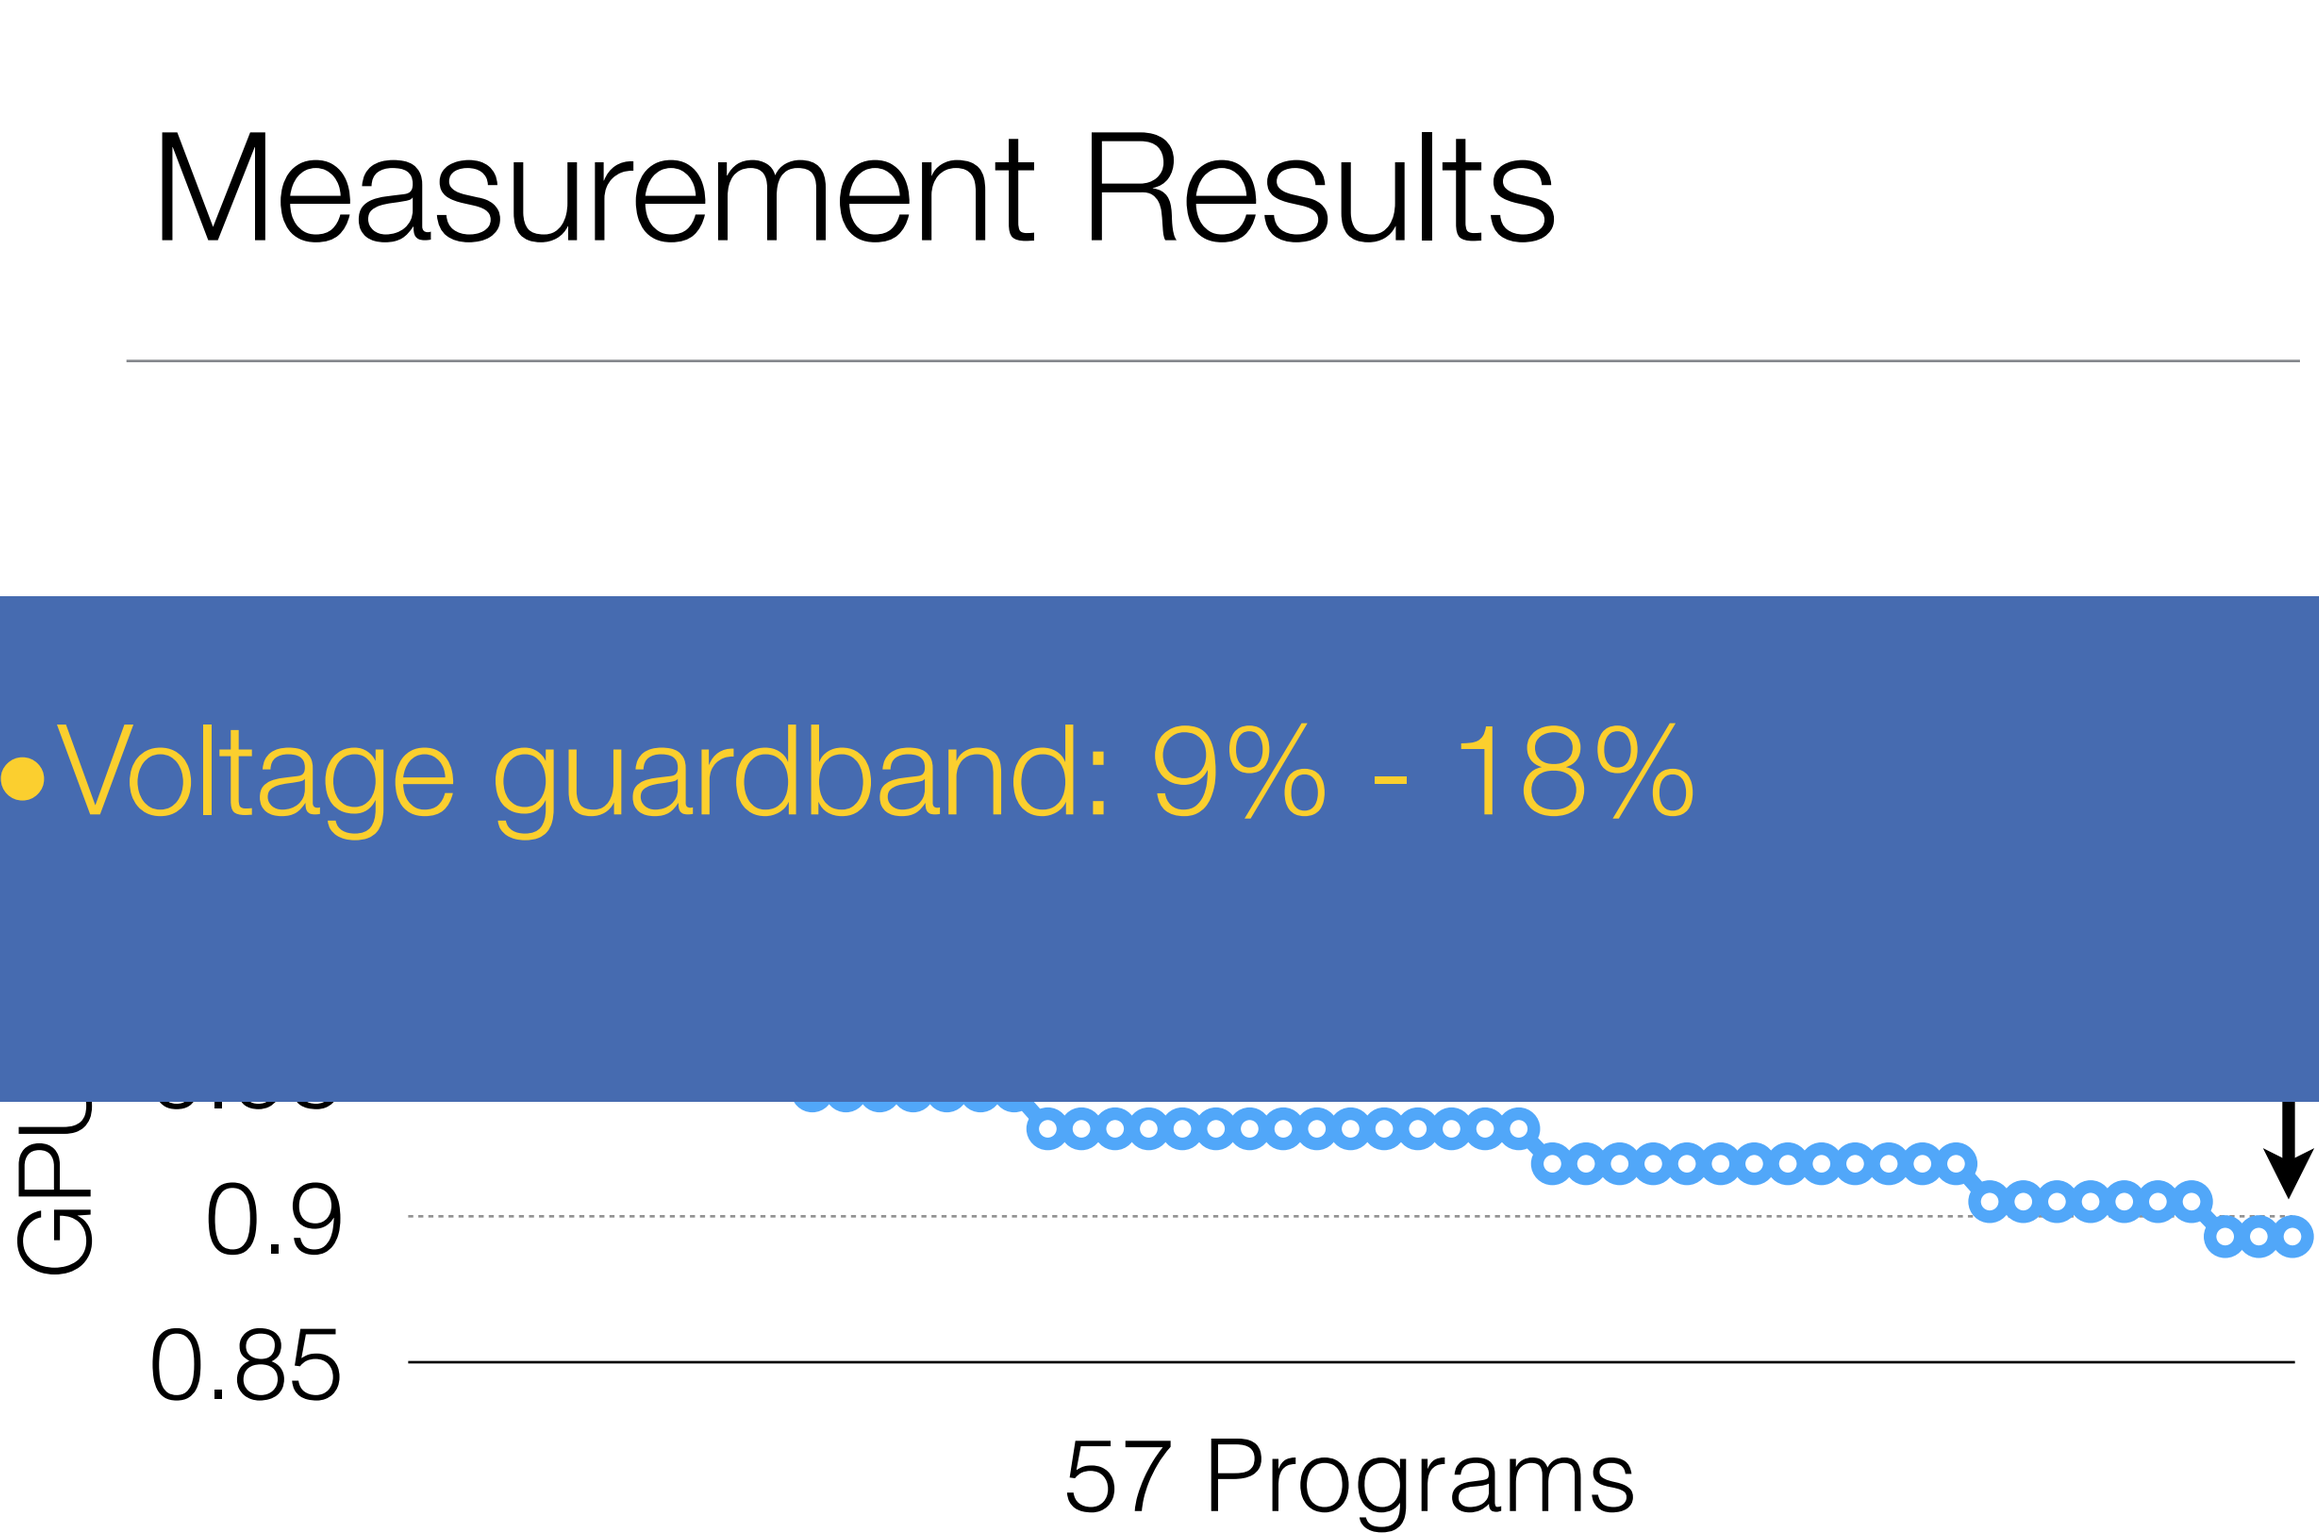

Measurement Results

- Voltage guardband: 9% - 18%
- Energy savings: 14% - 25% at the card level

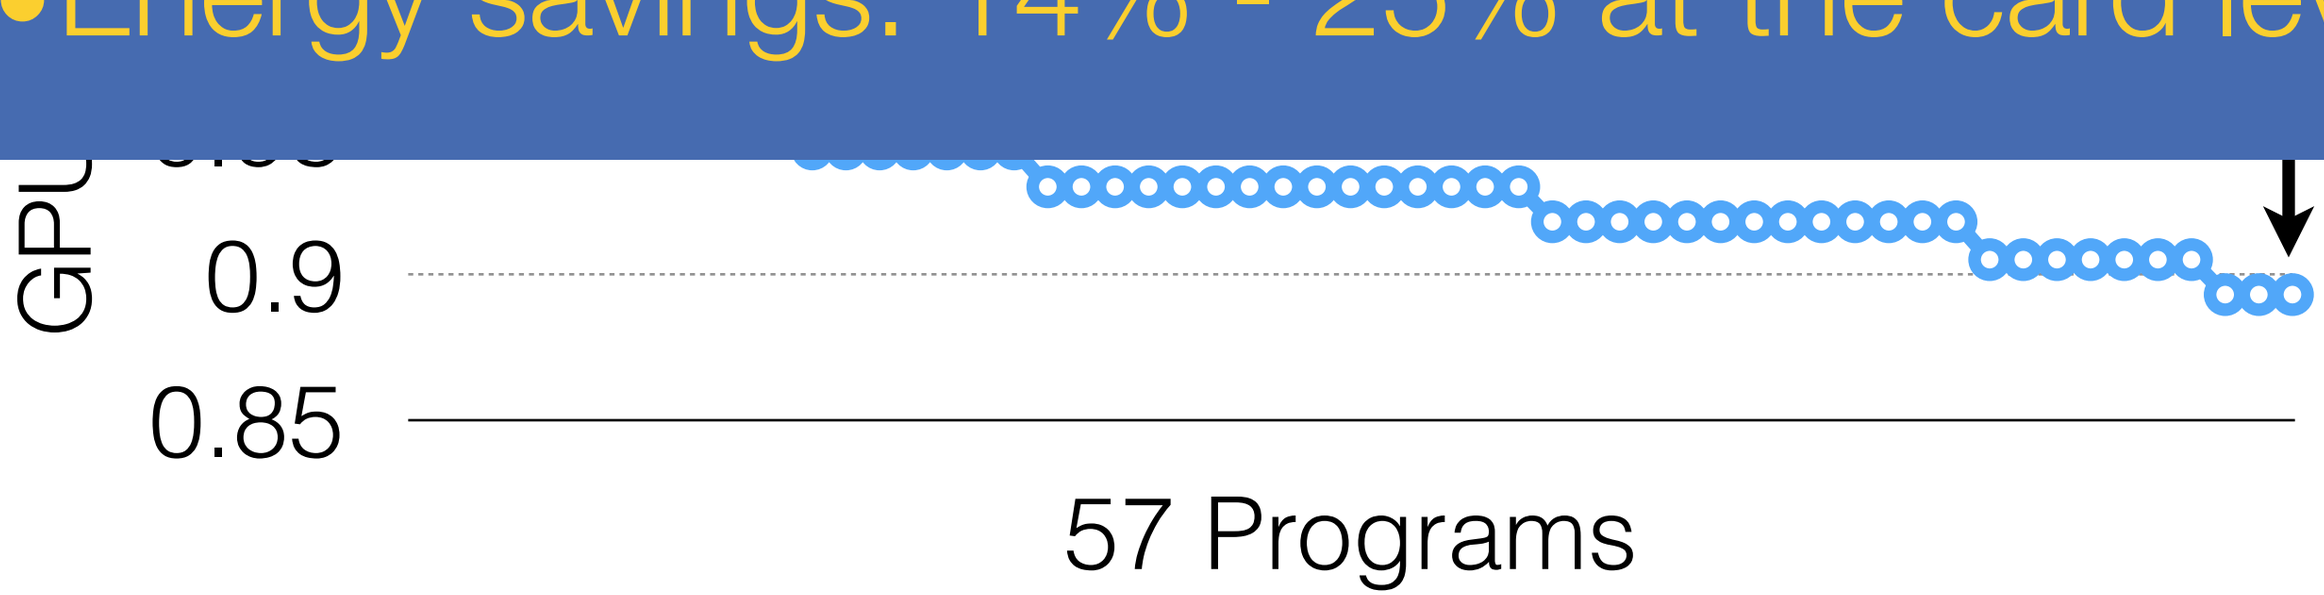

Measurement Results

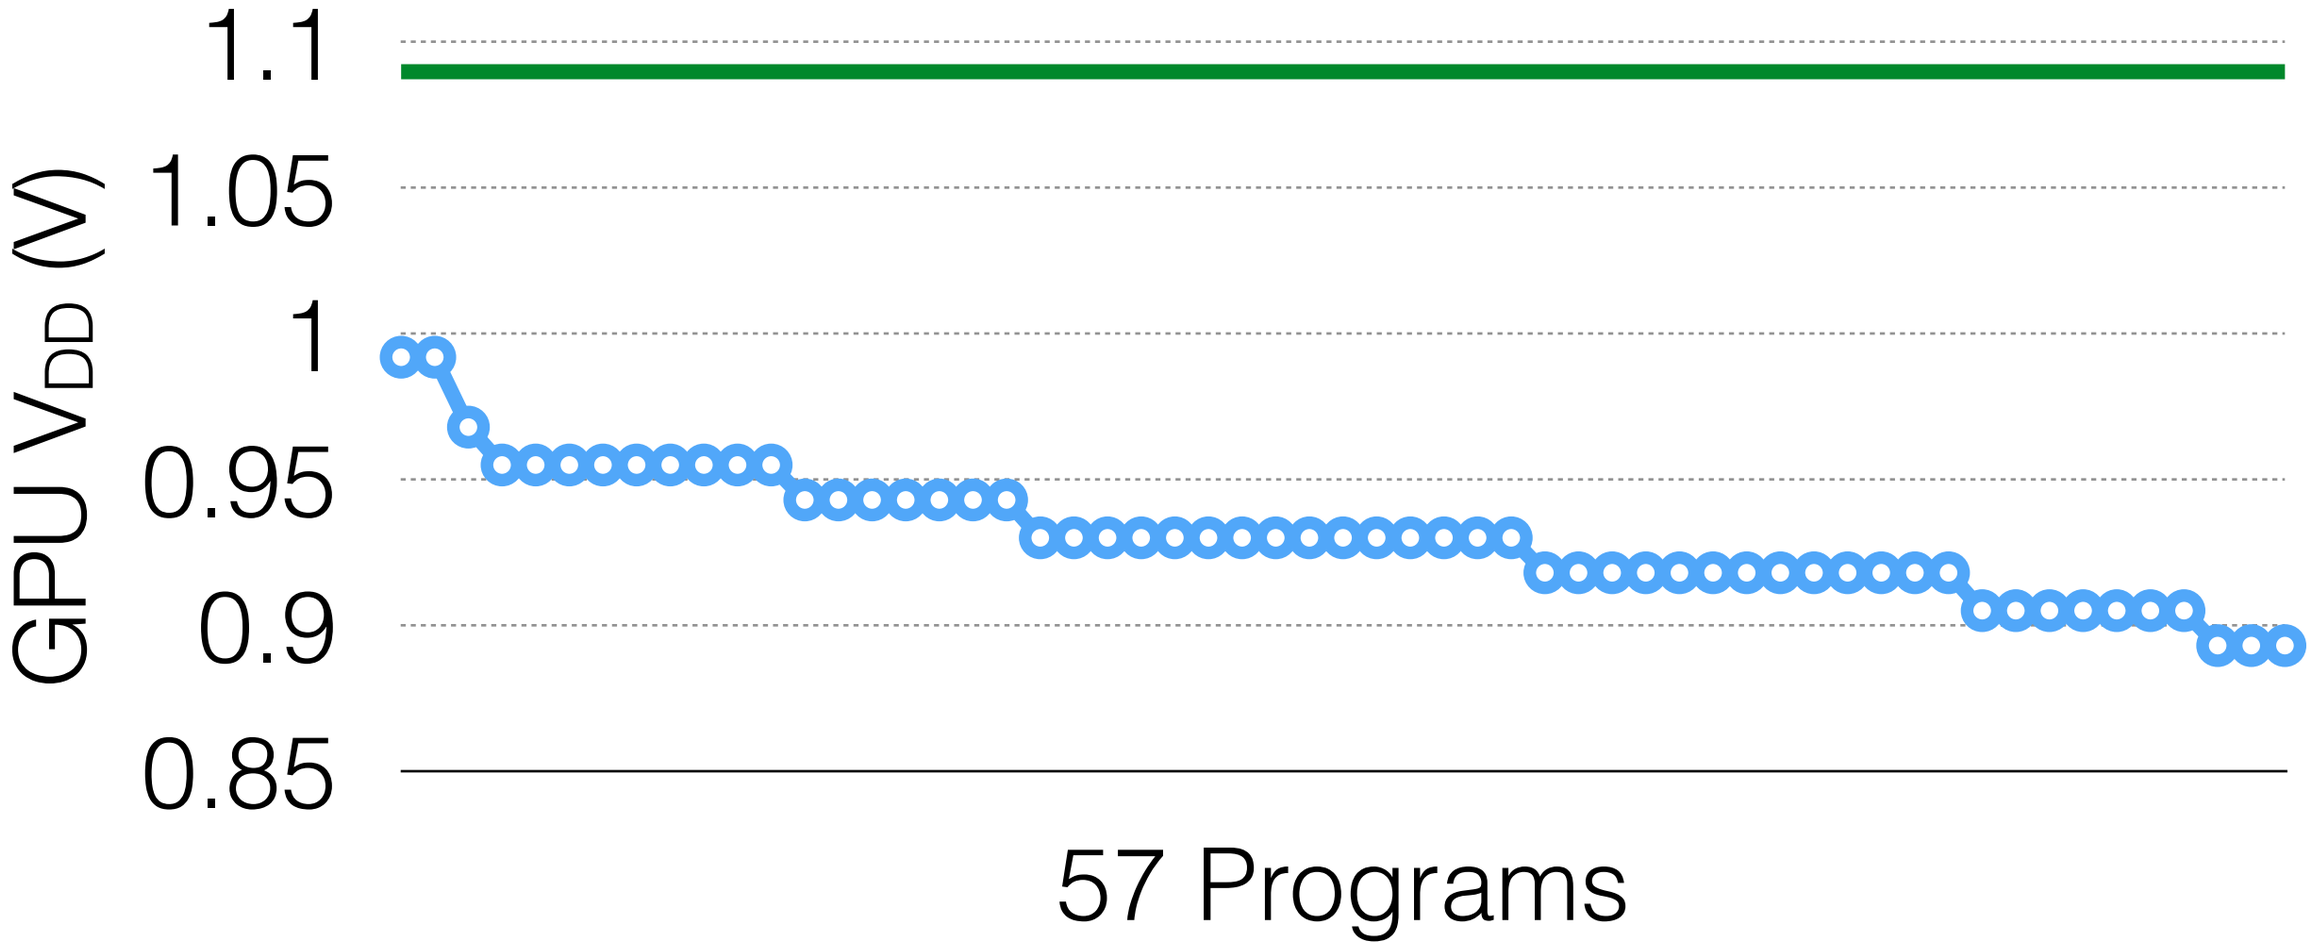

Measurement Results

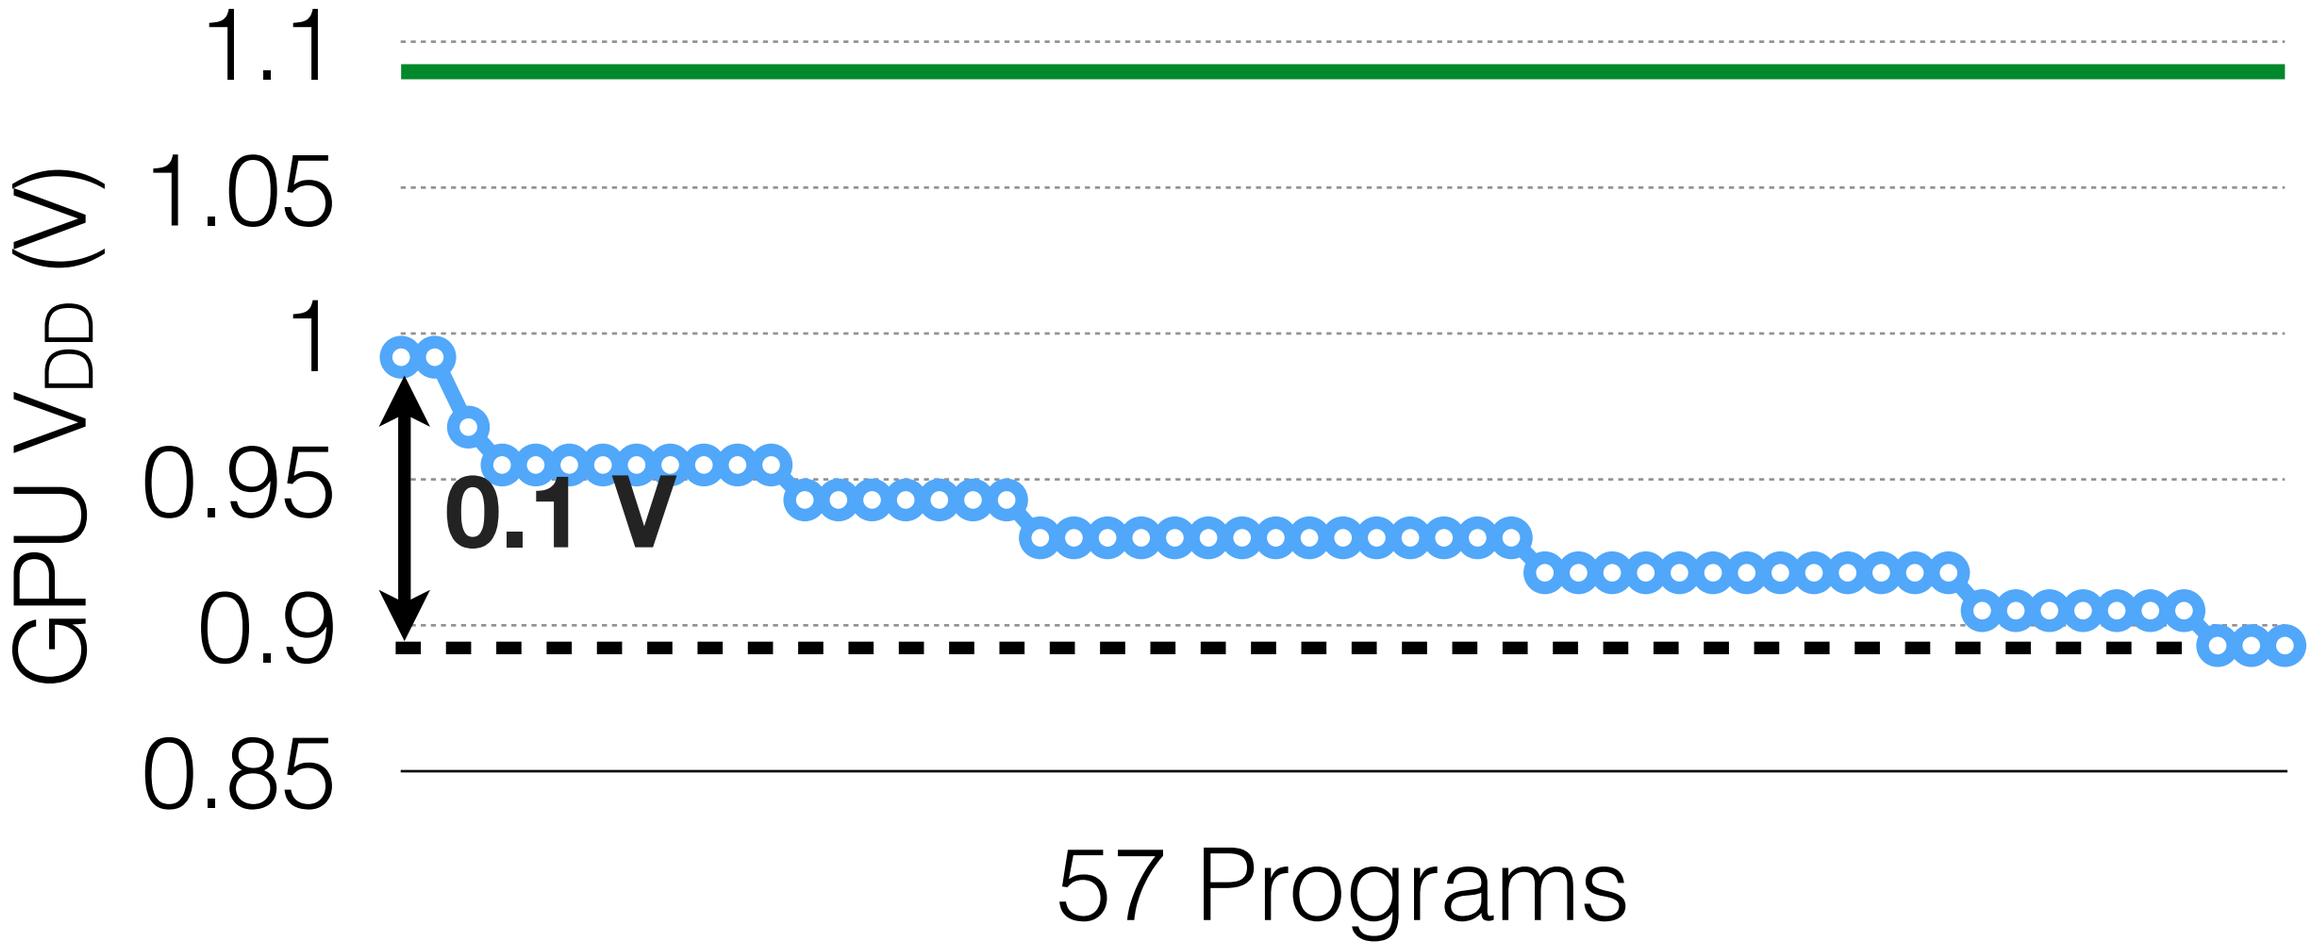

Measurement Results

V_{min} is program dependent

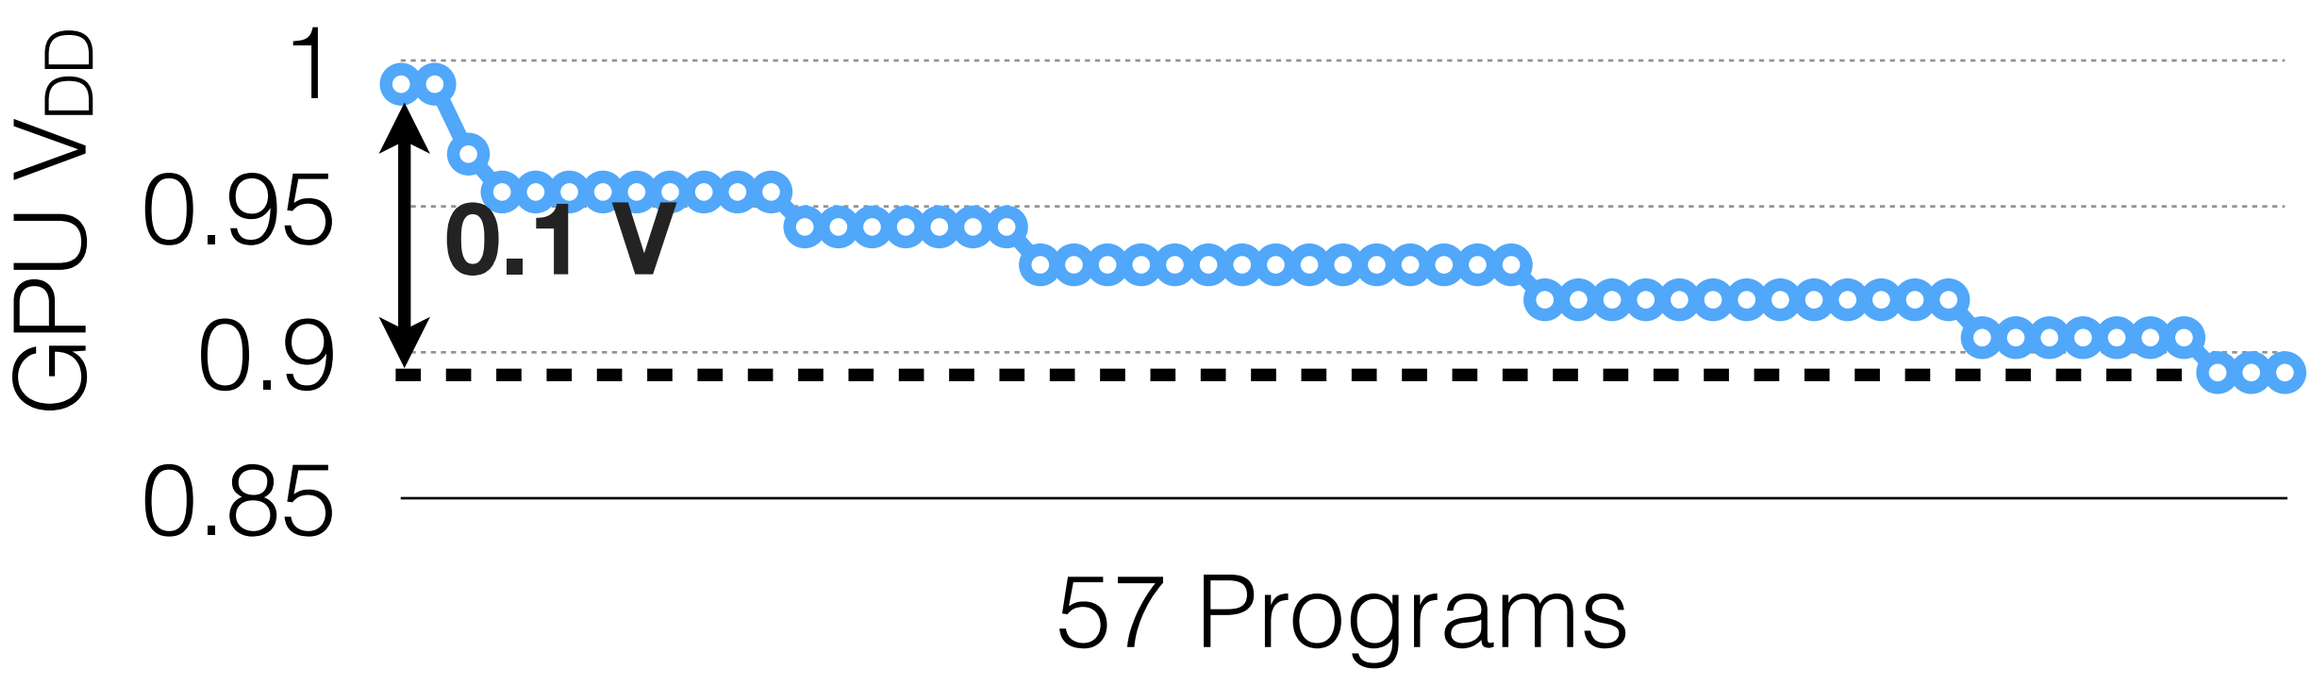

Executive Summary

- Guardband measurement

- Guardband analysis

- Guardband optimization

Voltage Guardband Analysis

Voltage Guardband Analysis

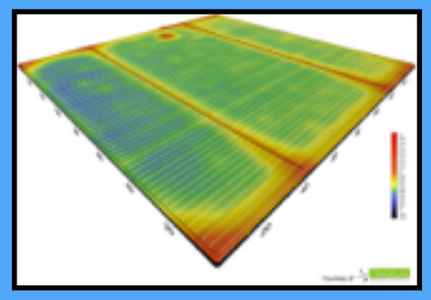

Process

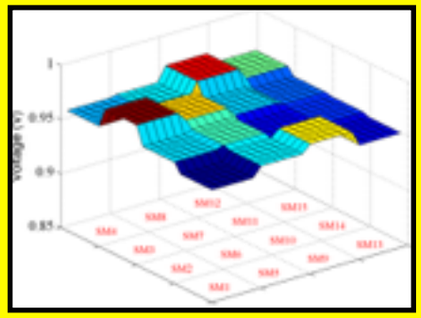

Voltage

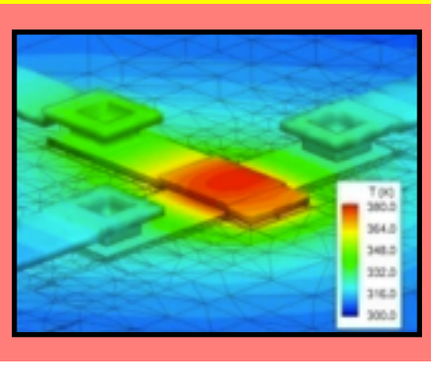

Temperature

Voltage Guardband Analysis

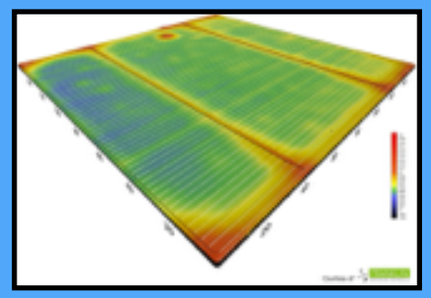

Process

?%

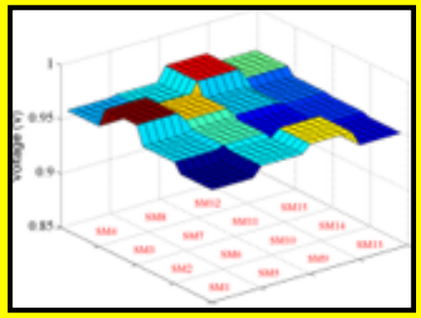

Voltage

?%

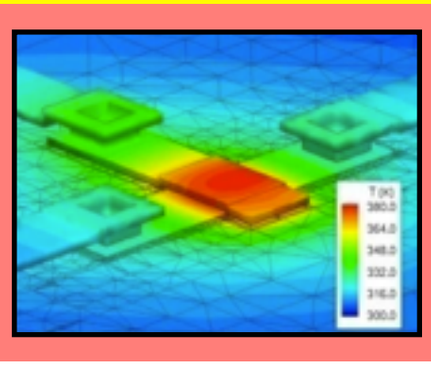

Temperature

?%

Voltage Guardband Analysis

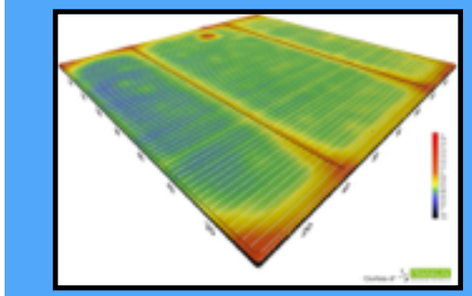

Process

?%

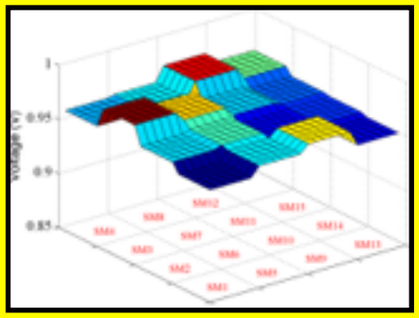

Voltage

?%

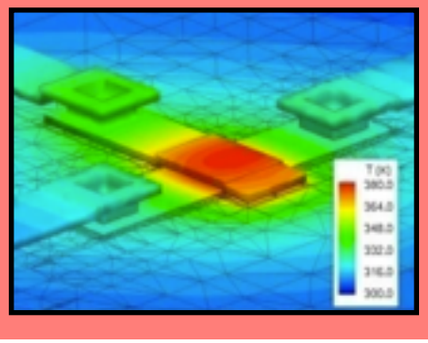

Temperature

?%

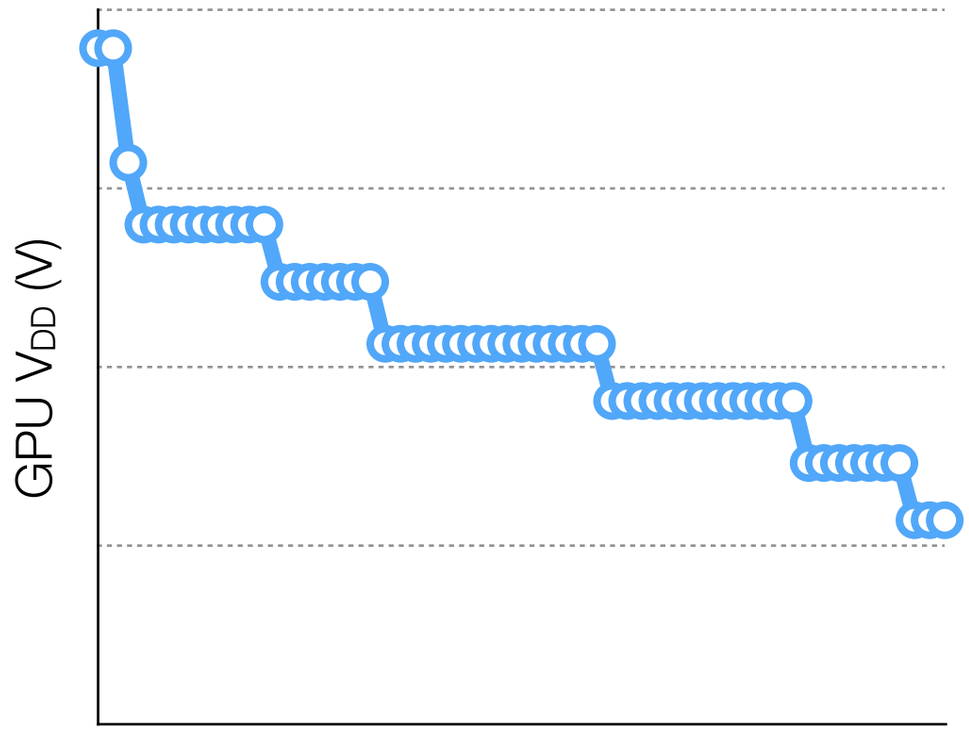

Voltage Guardband Analysis

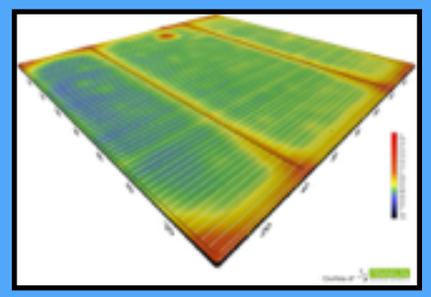

Process

?%

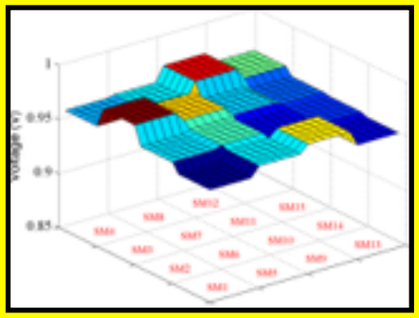

Voltage

?%

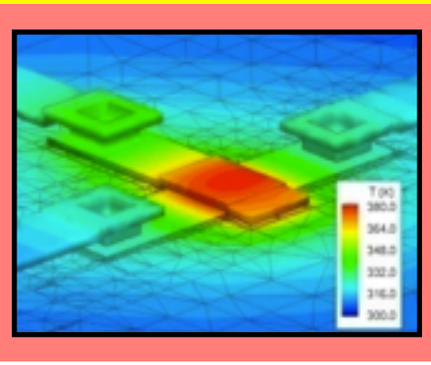

Temperature

?%

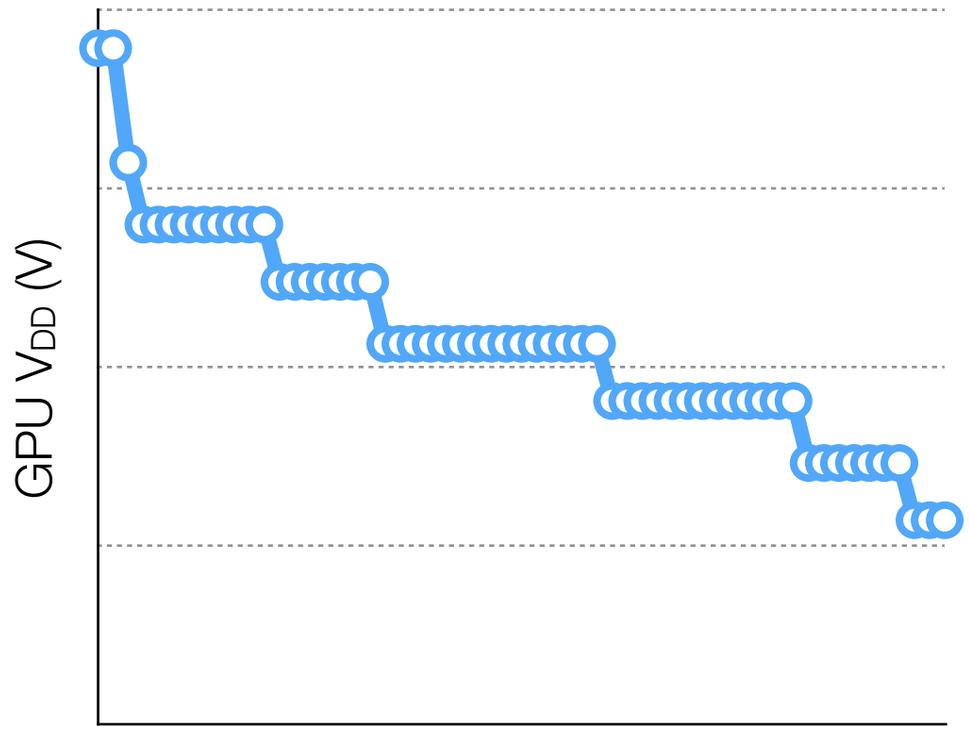

Voltage Guardband Analysis

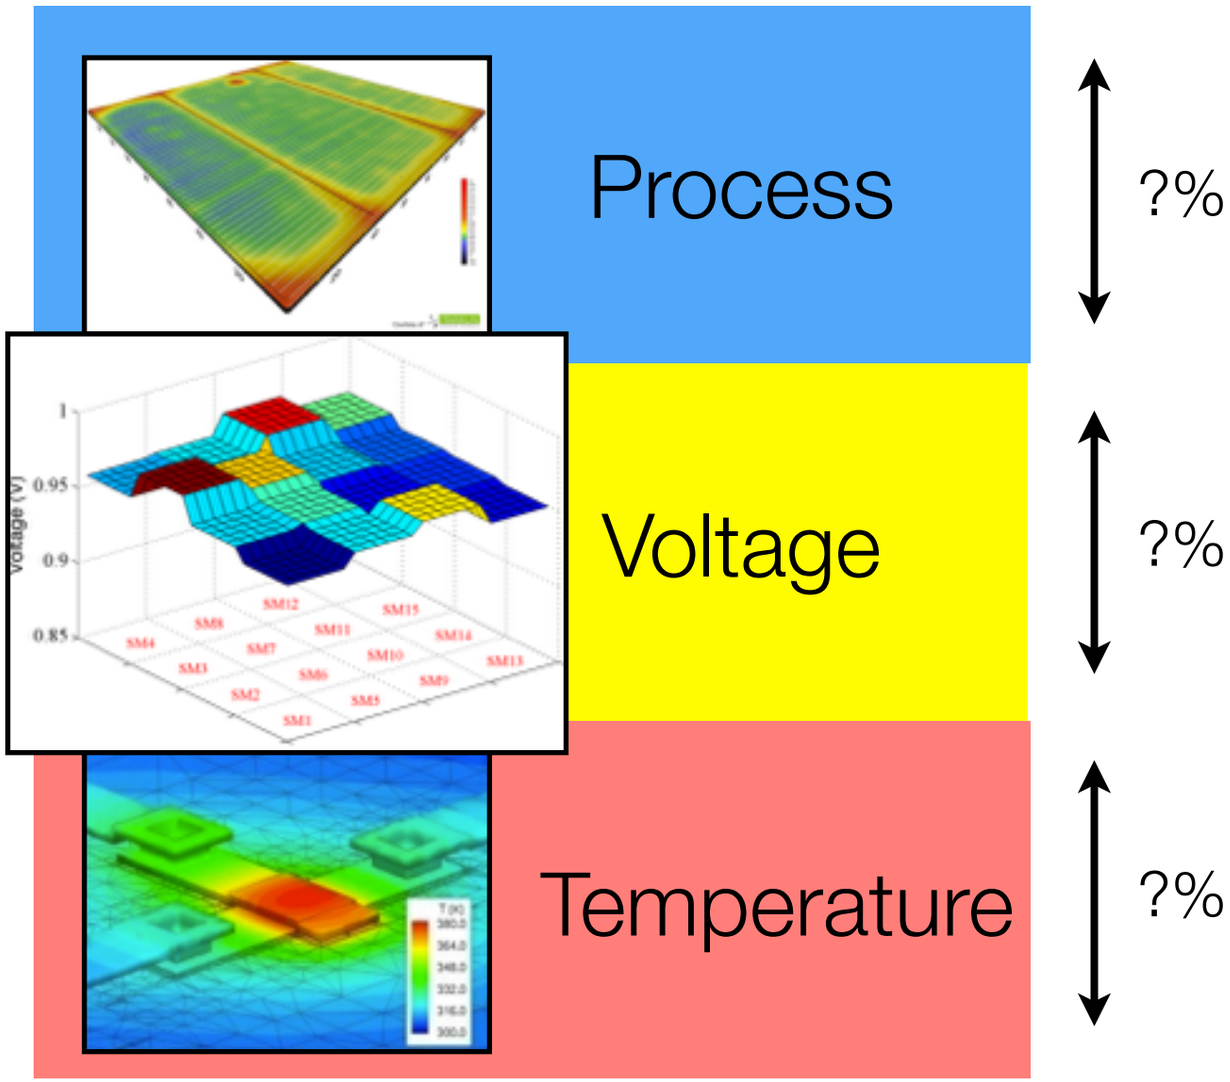

Process Variation Impact

Process Variation Impact

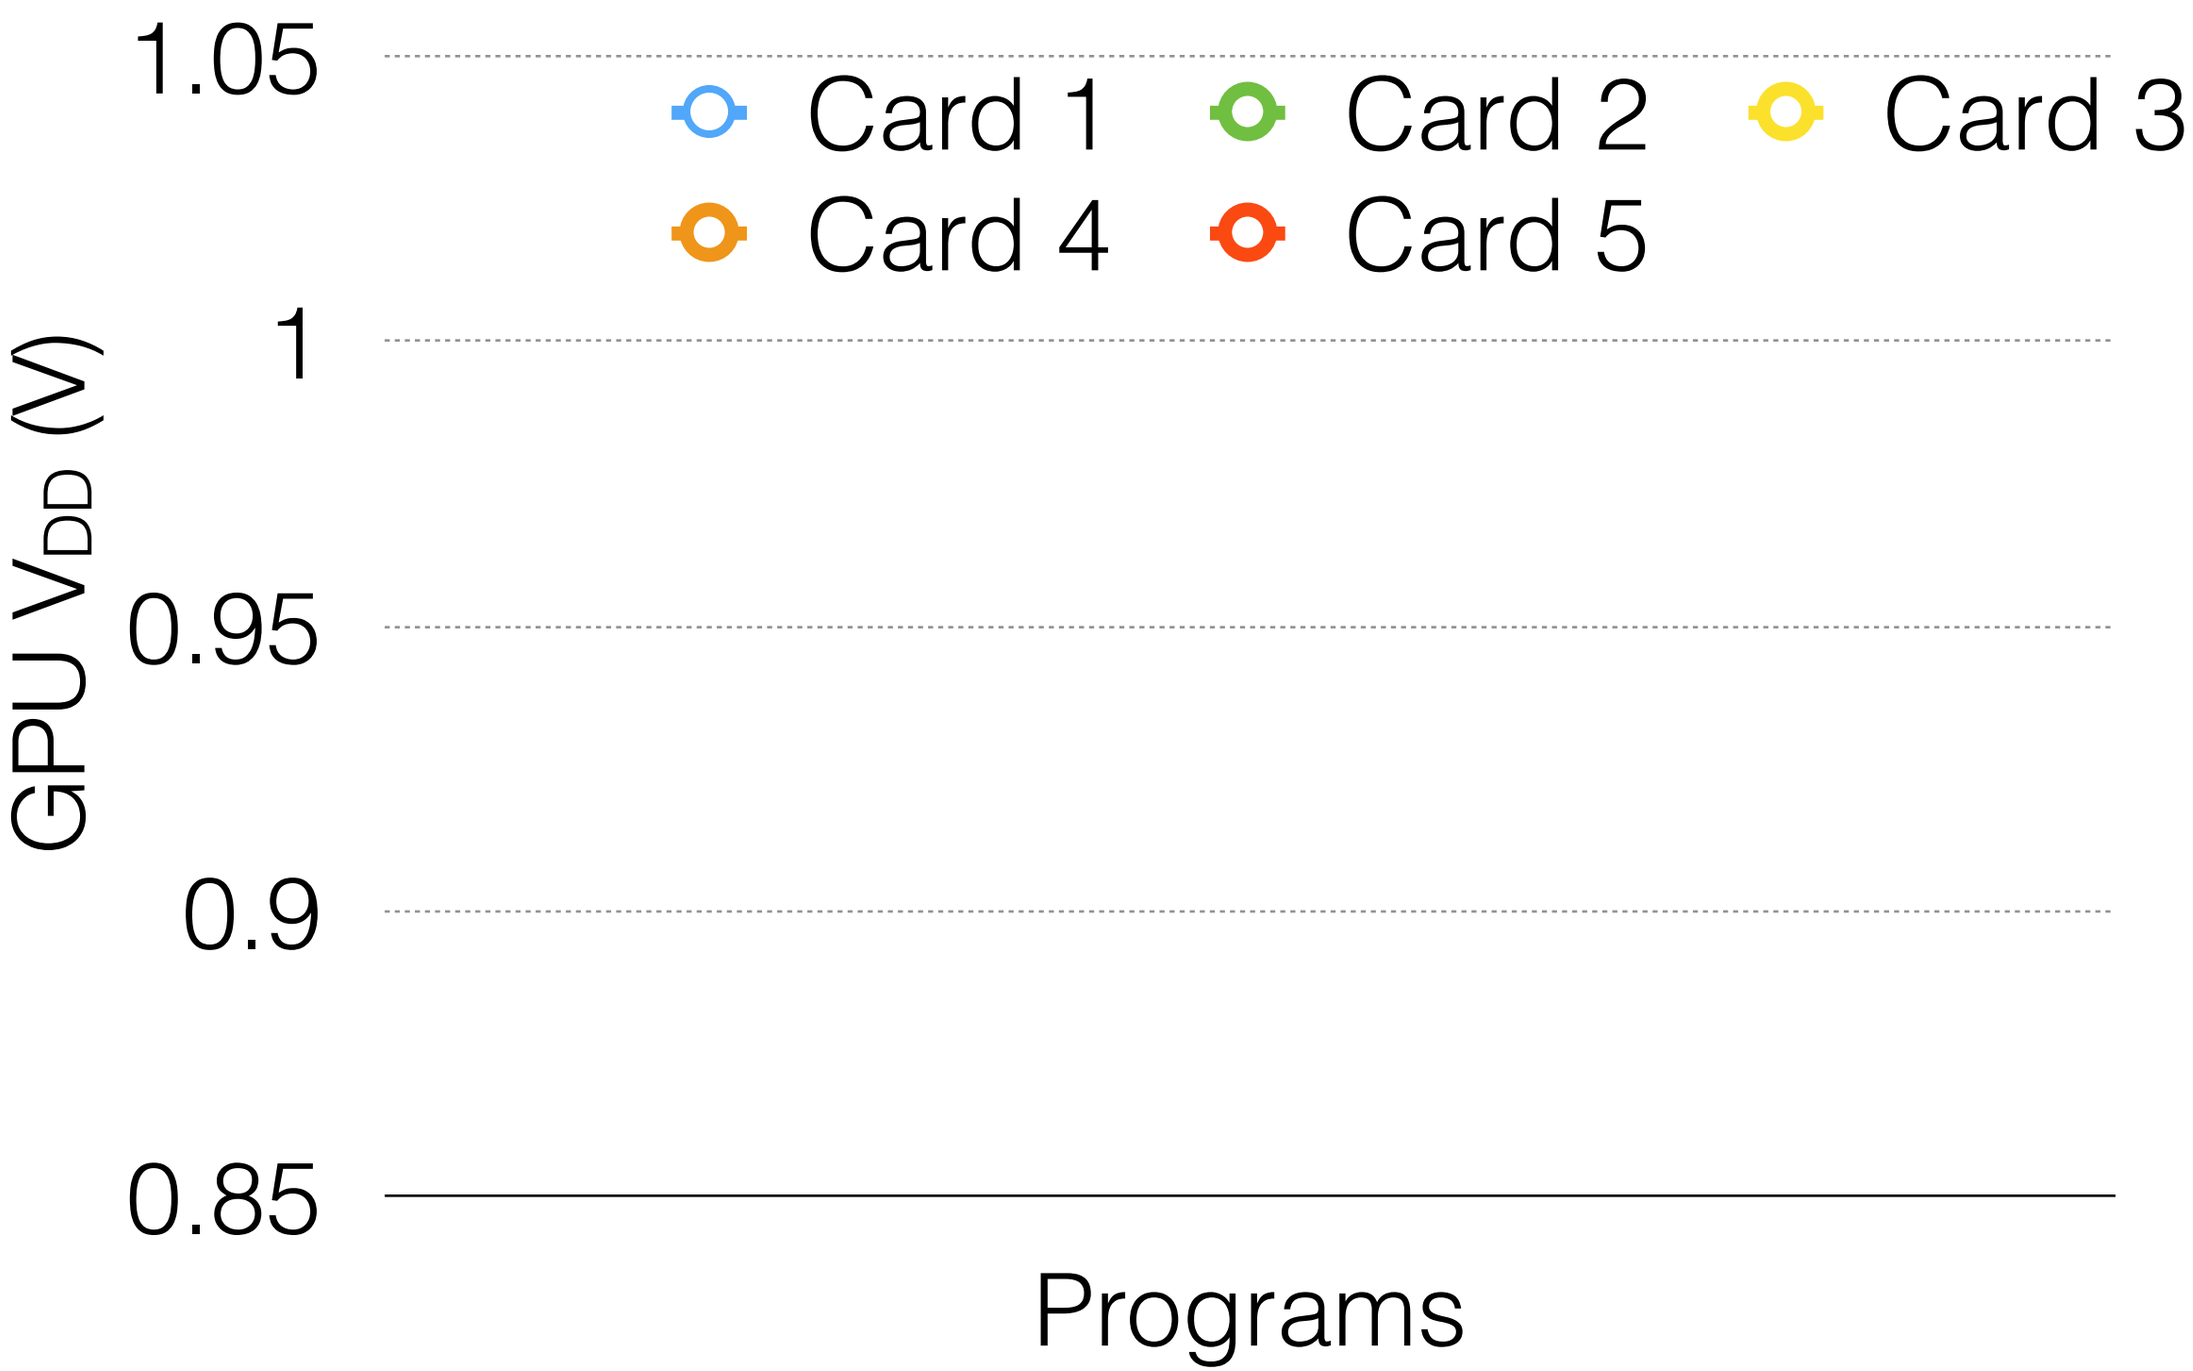

Process Variation Impact

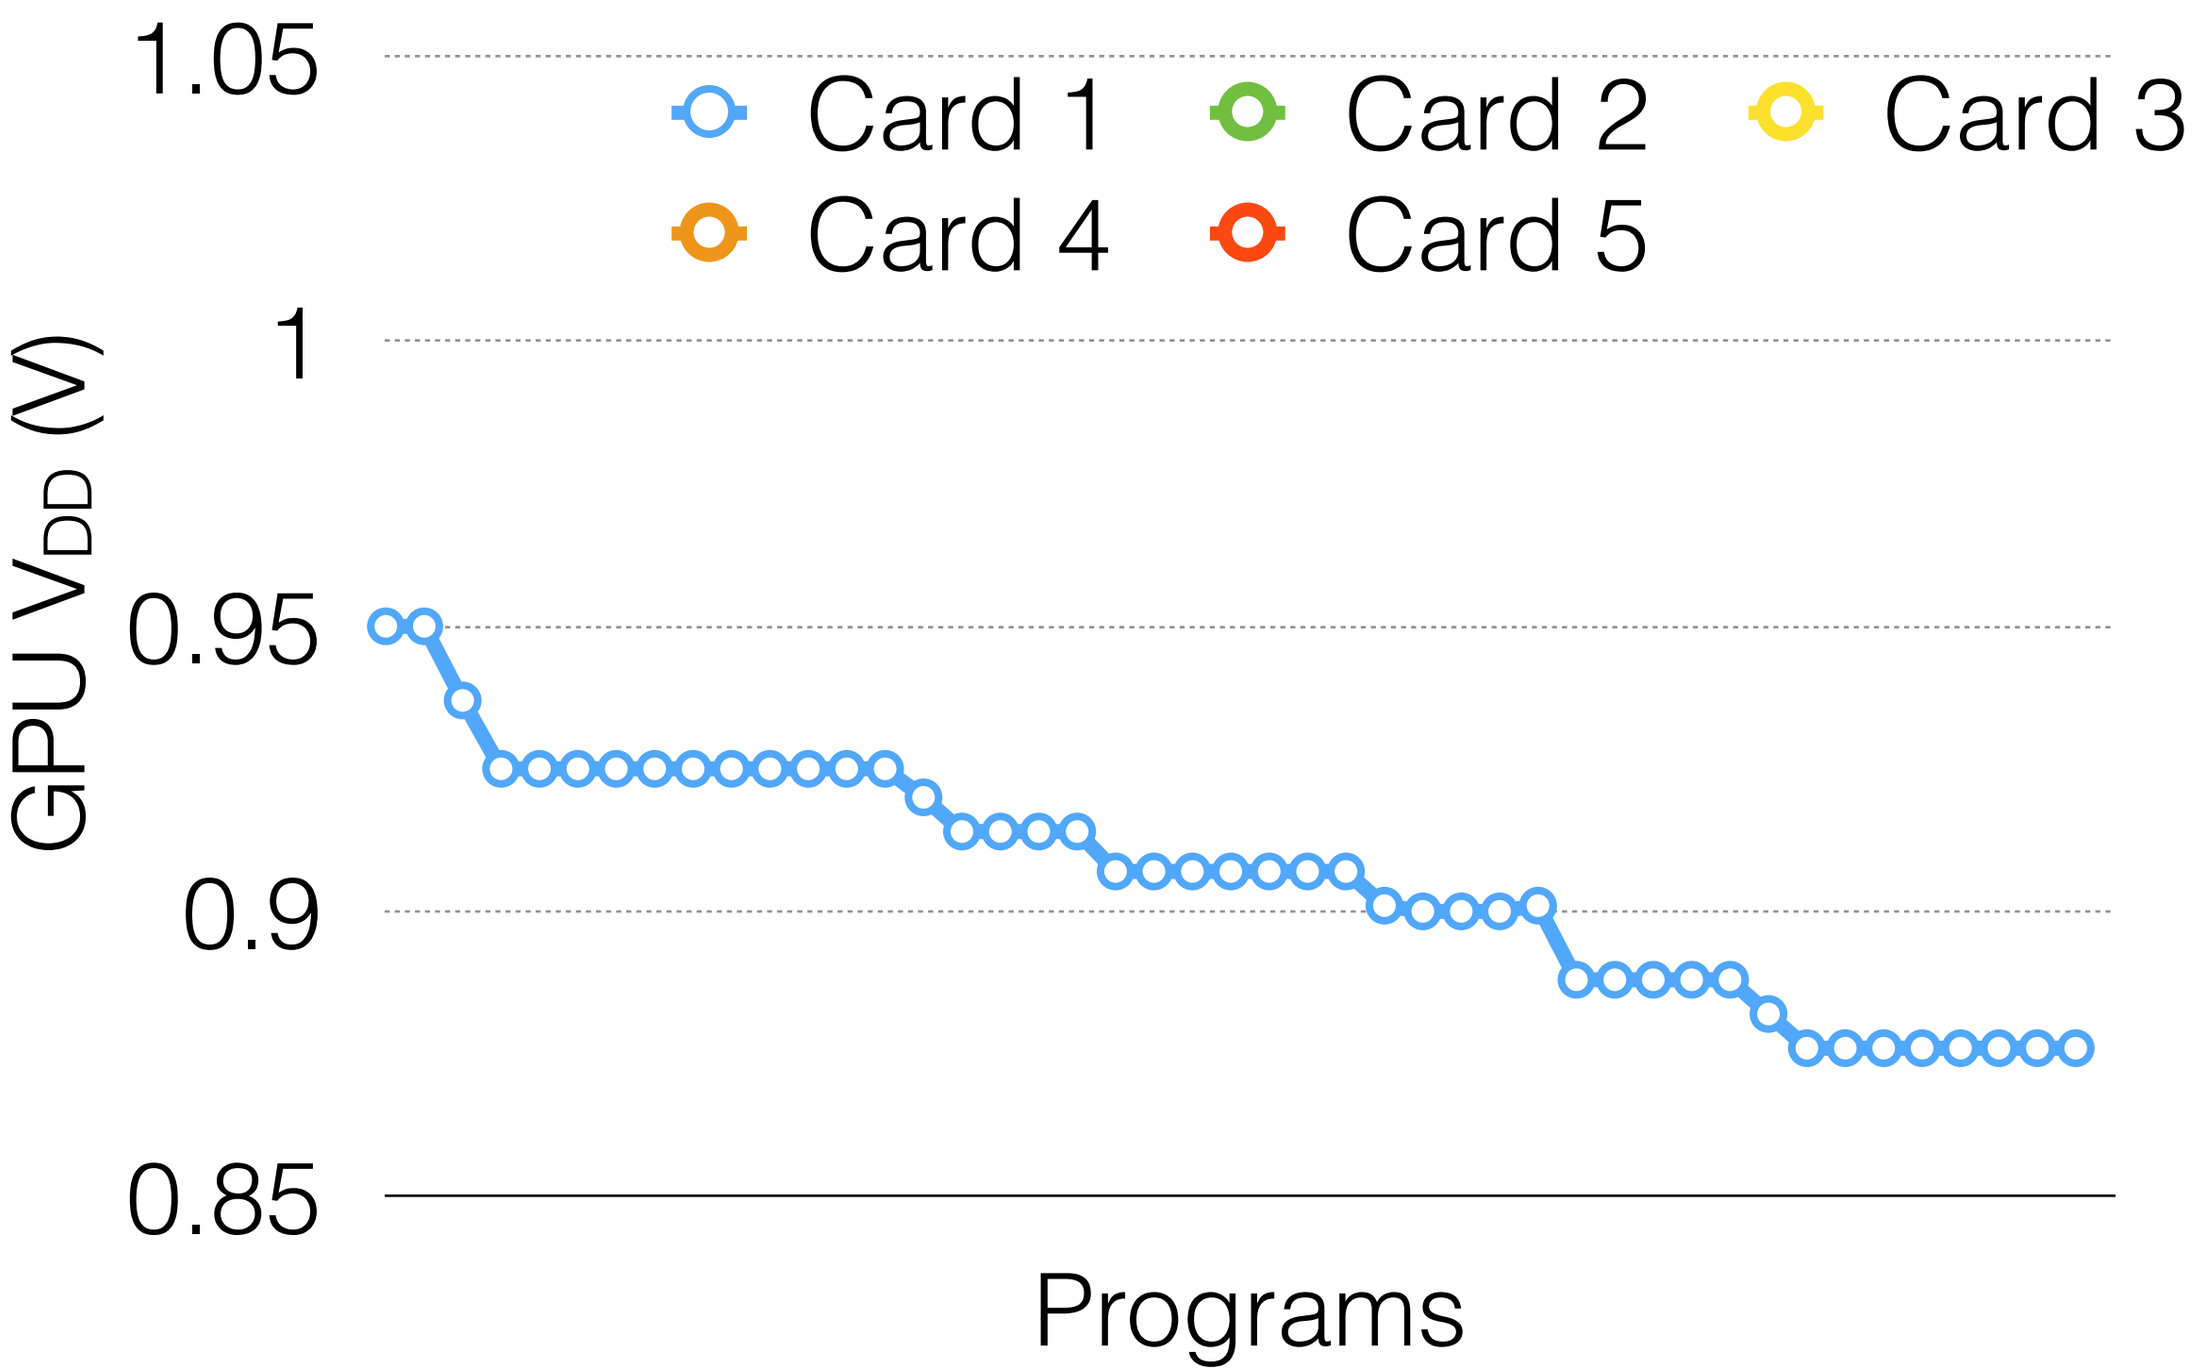

Process Variation Impact

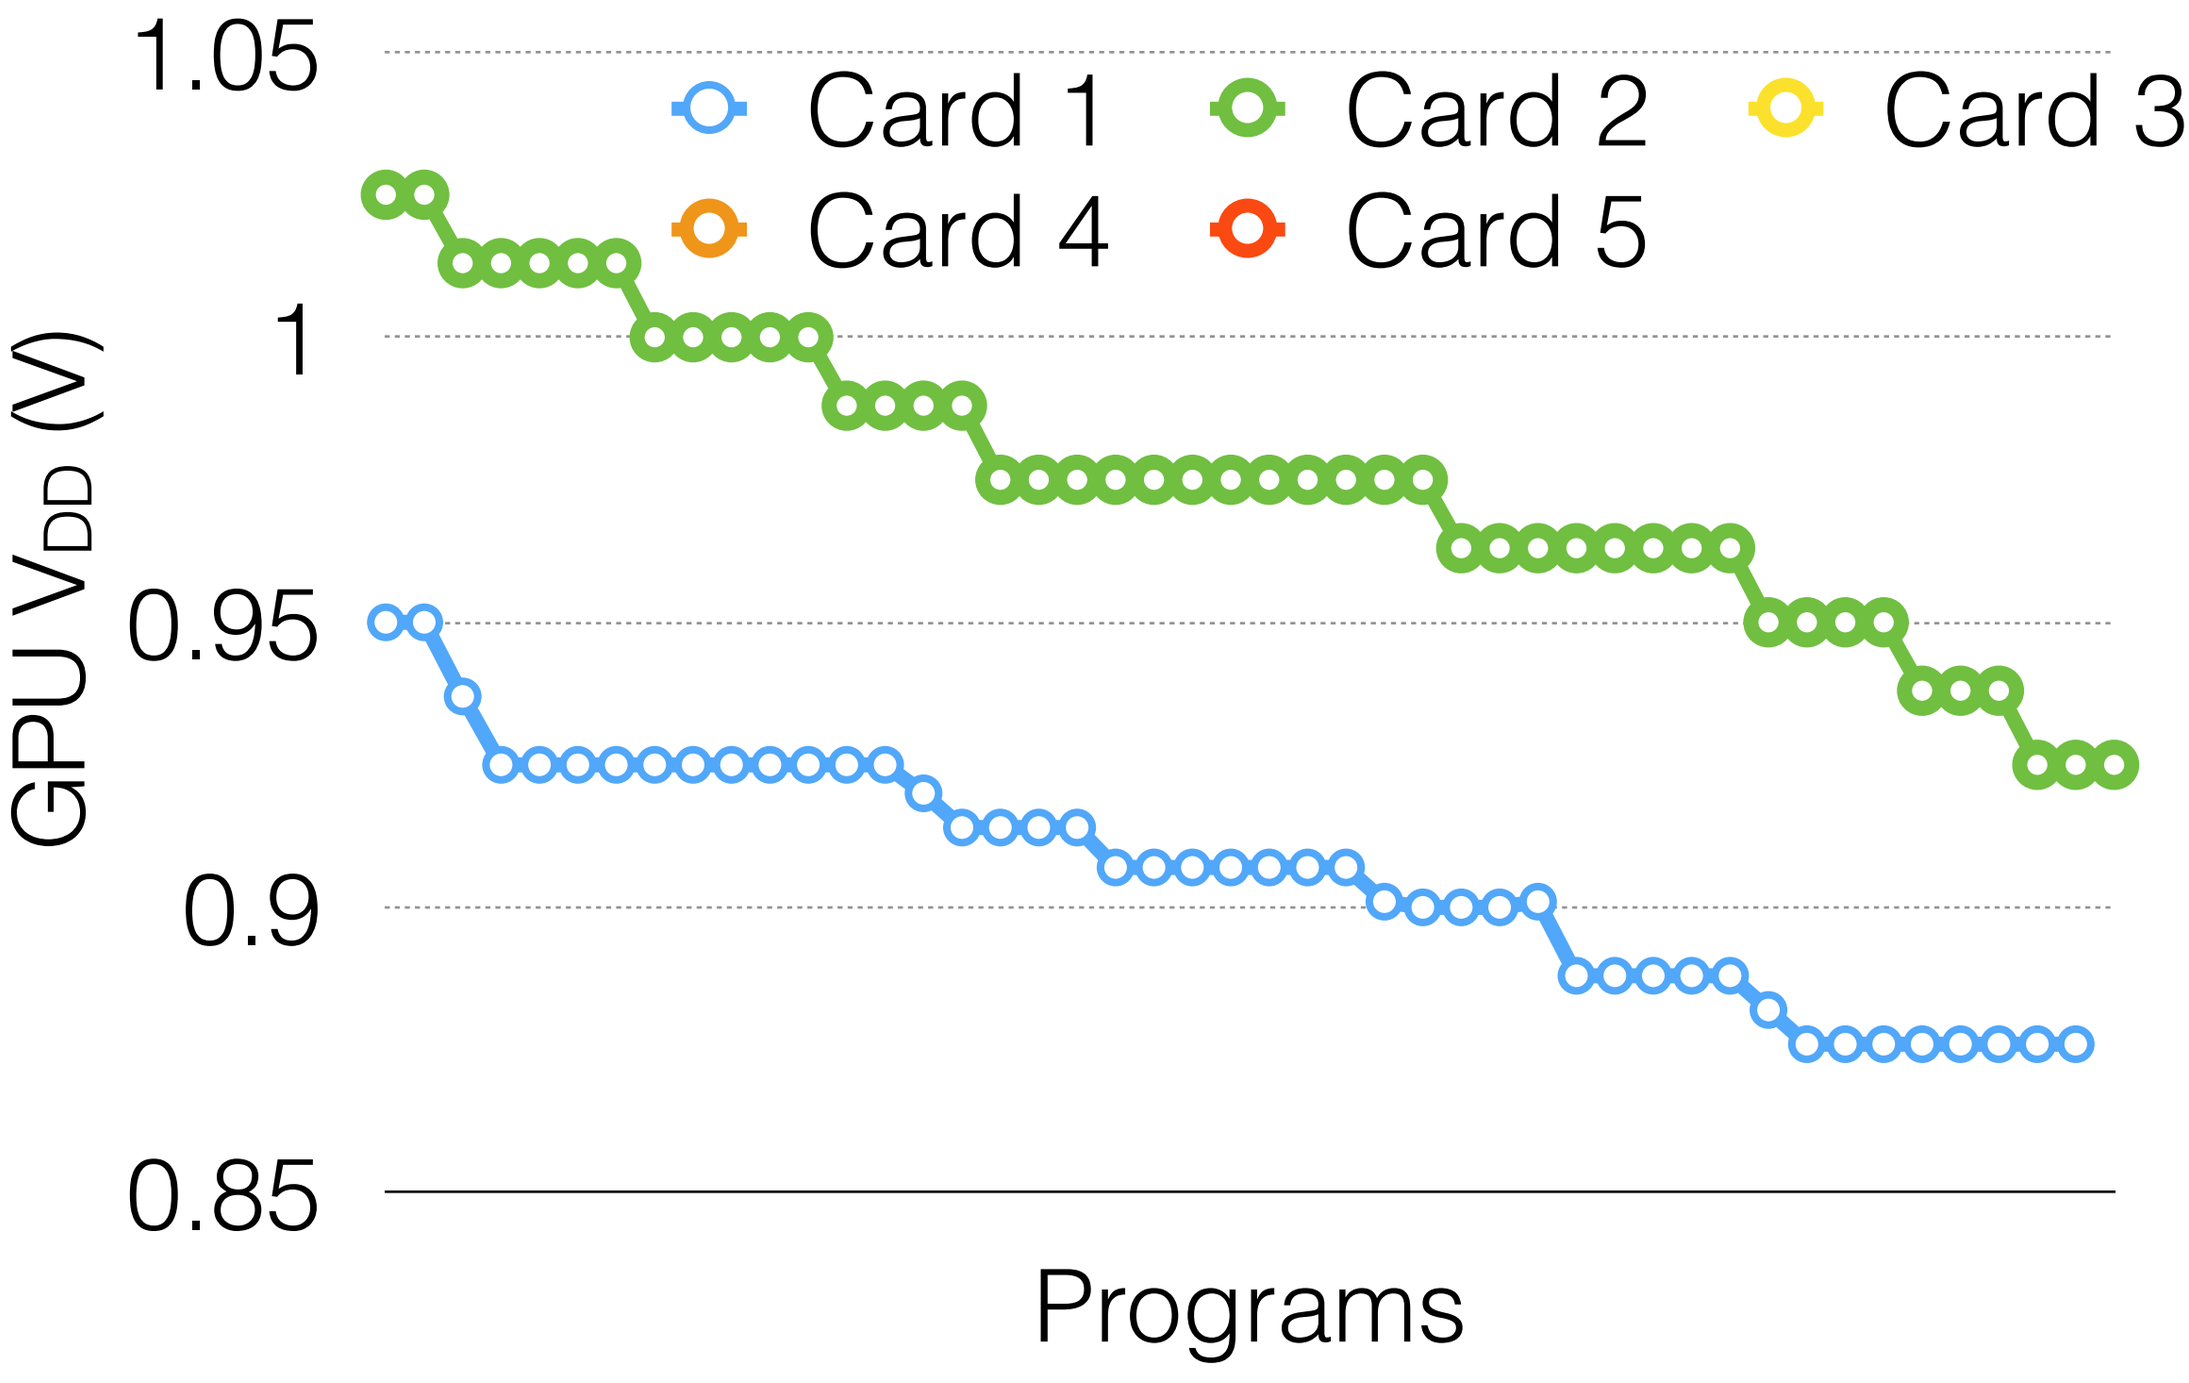

Process Variation Impact

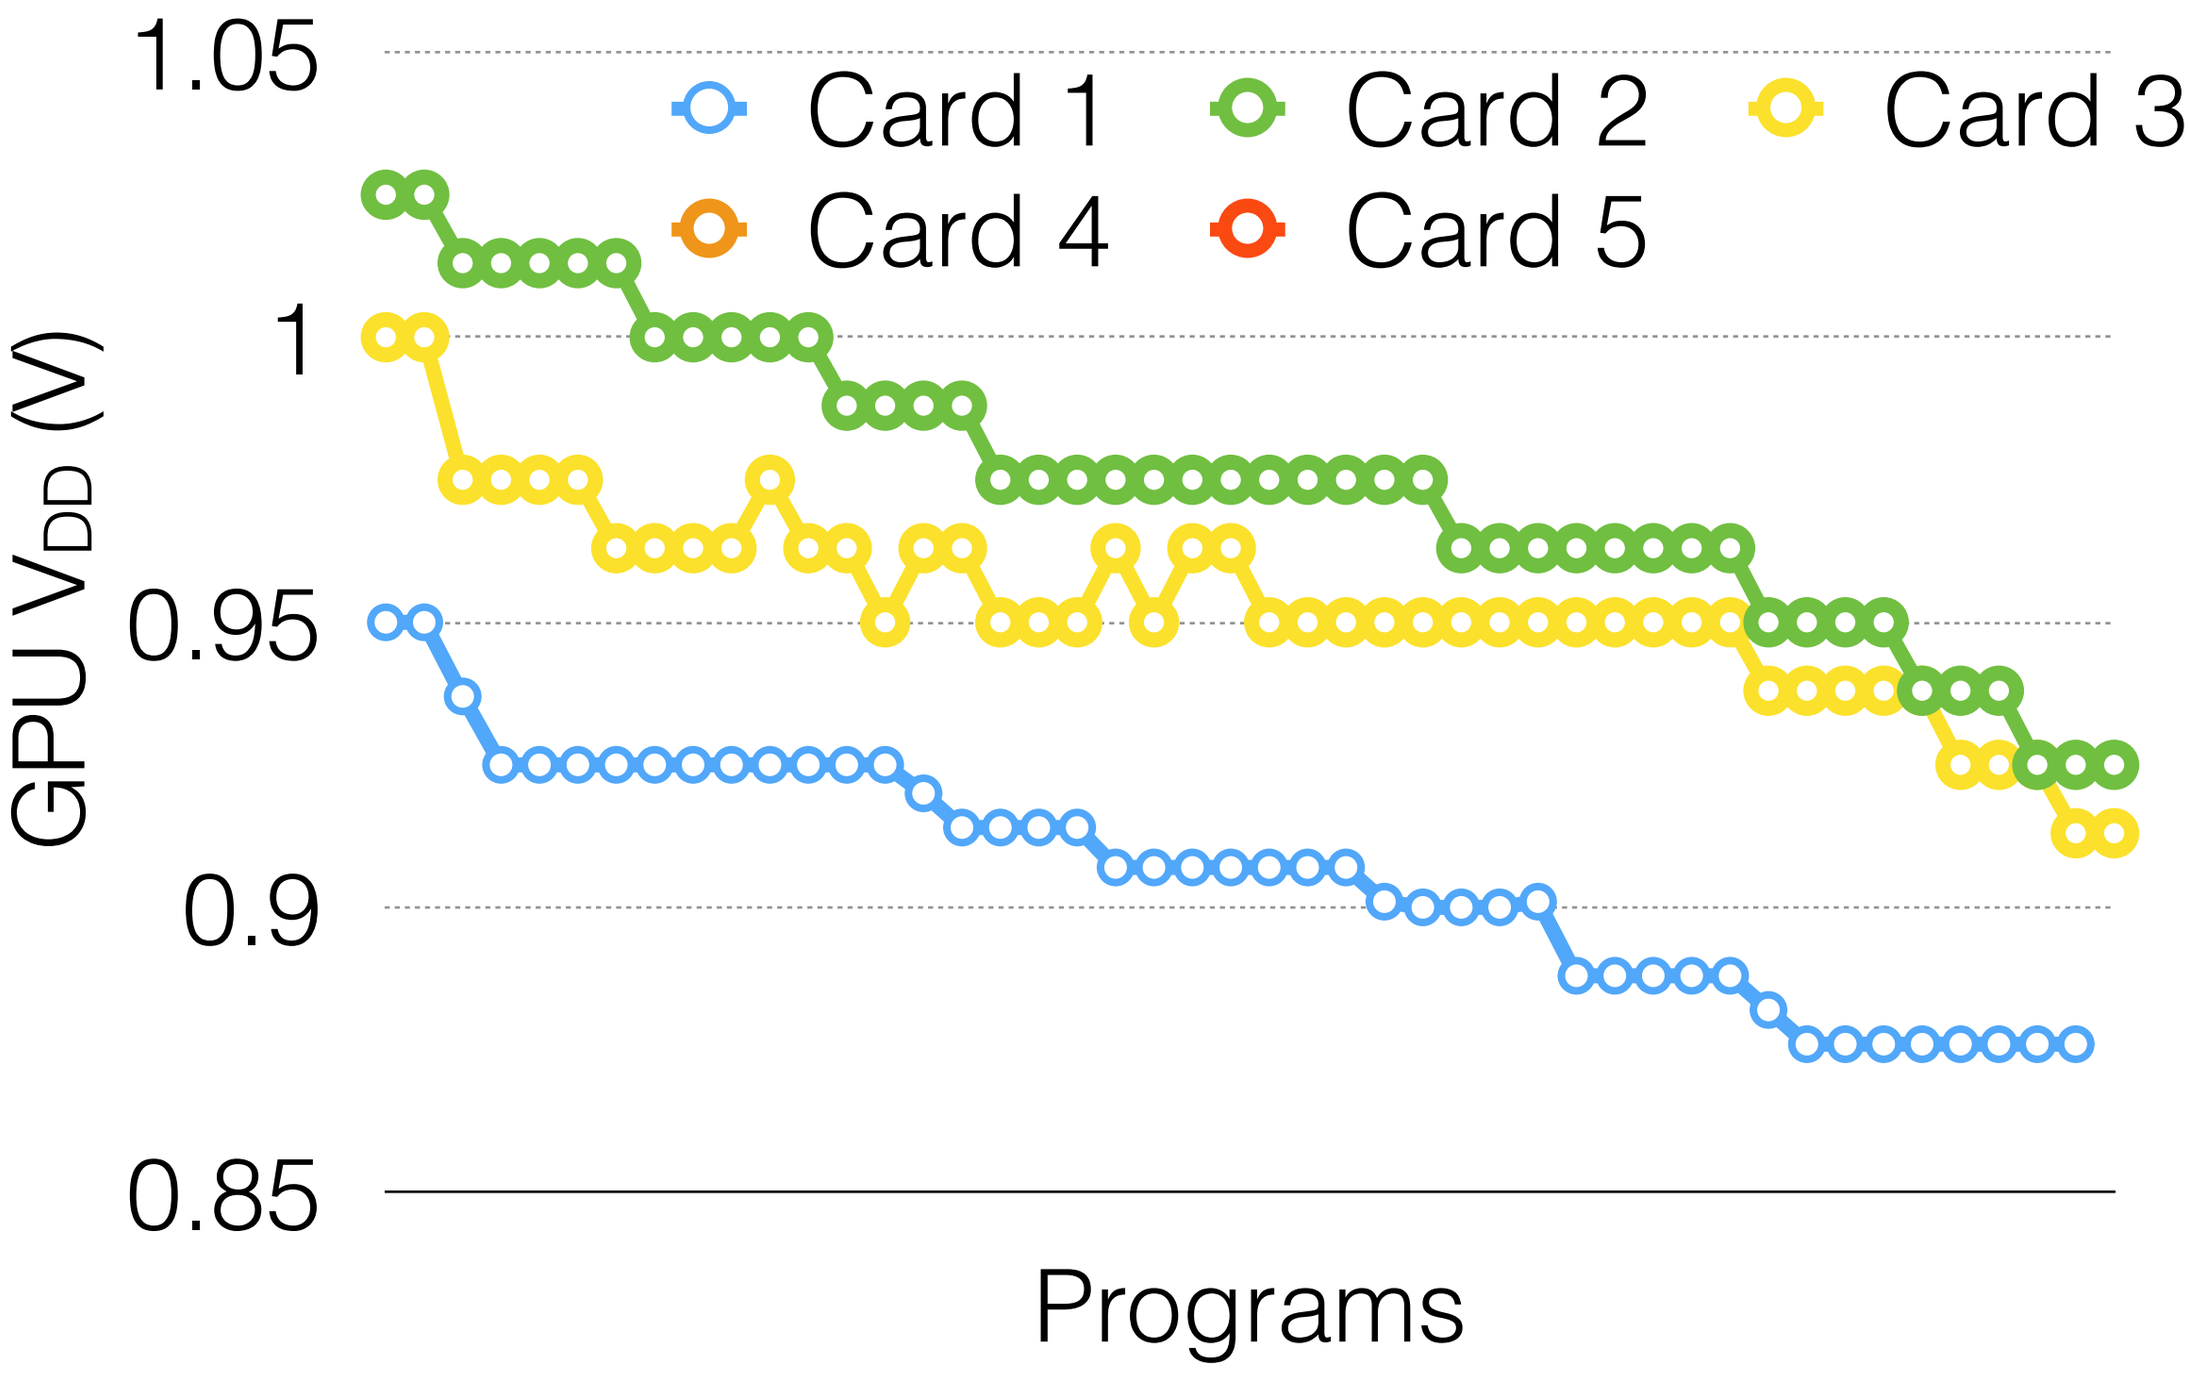

Process Variation Impact

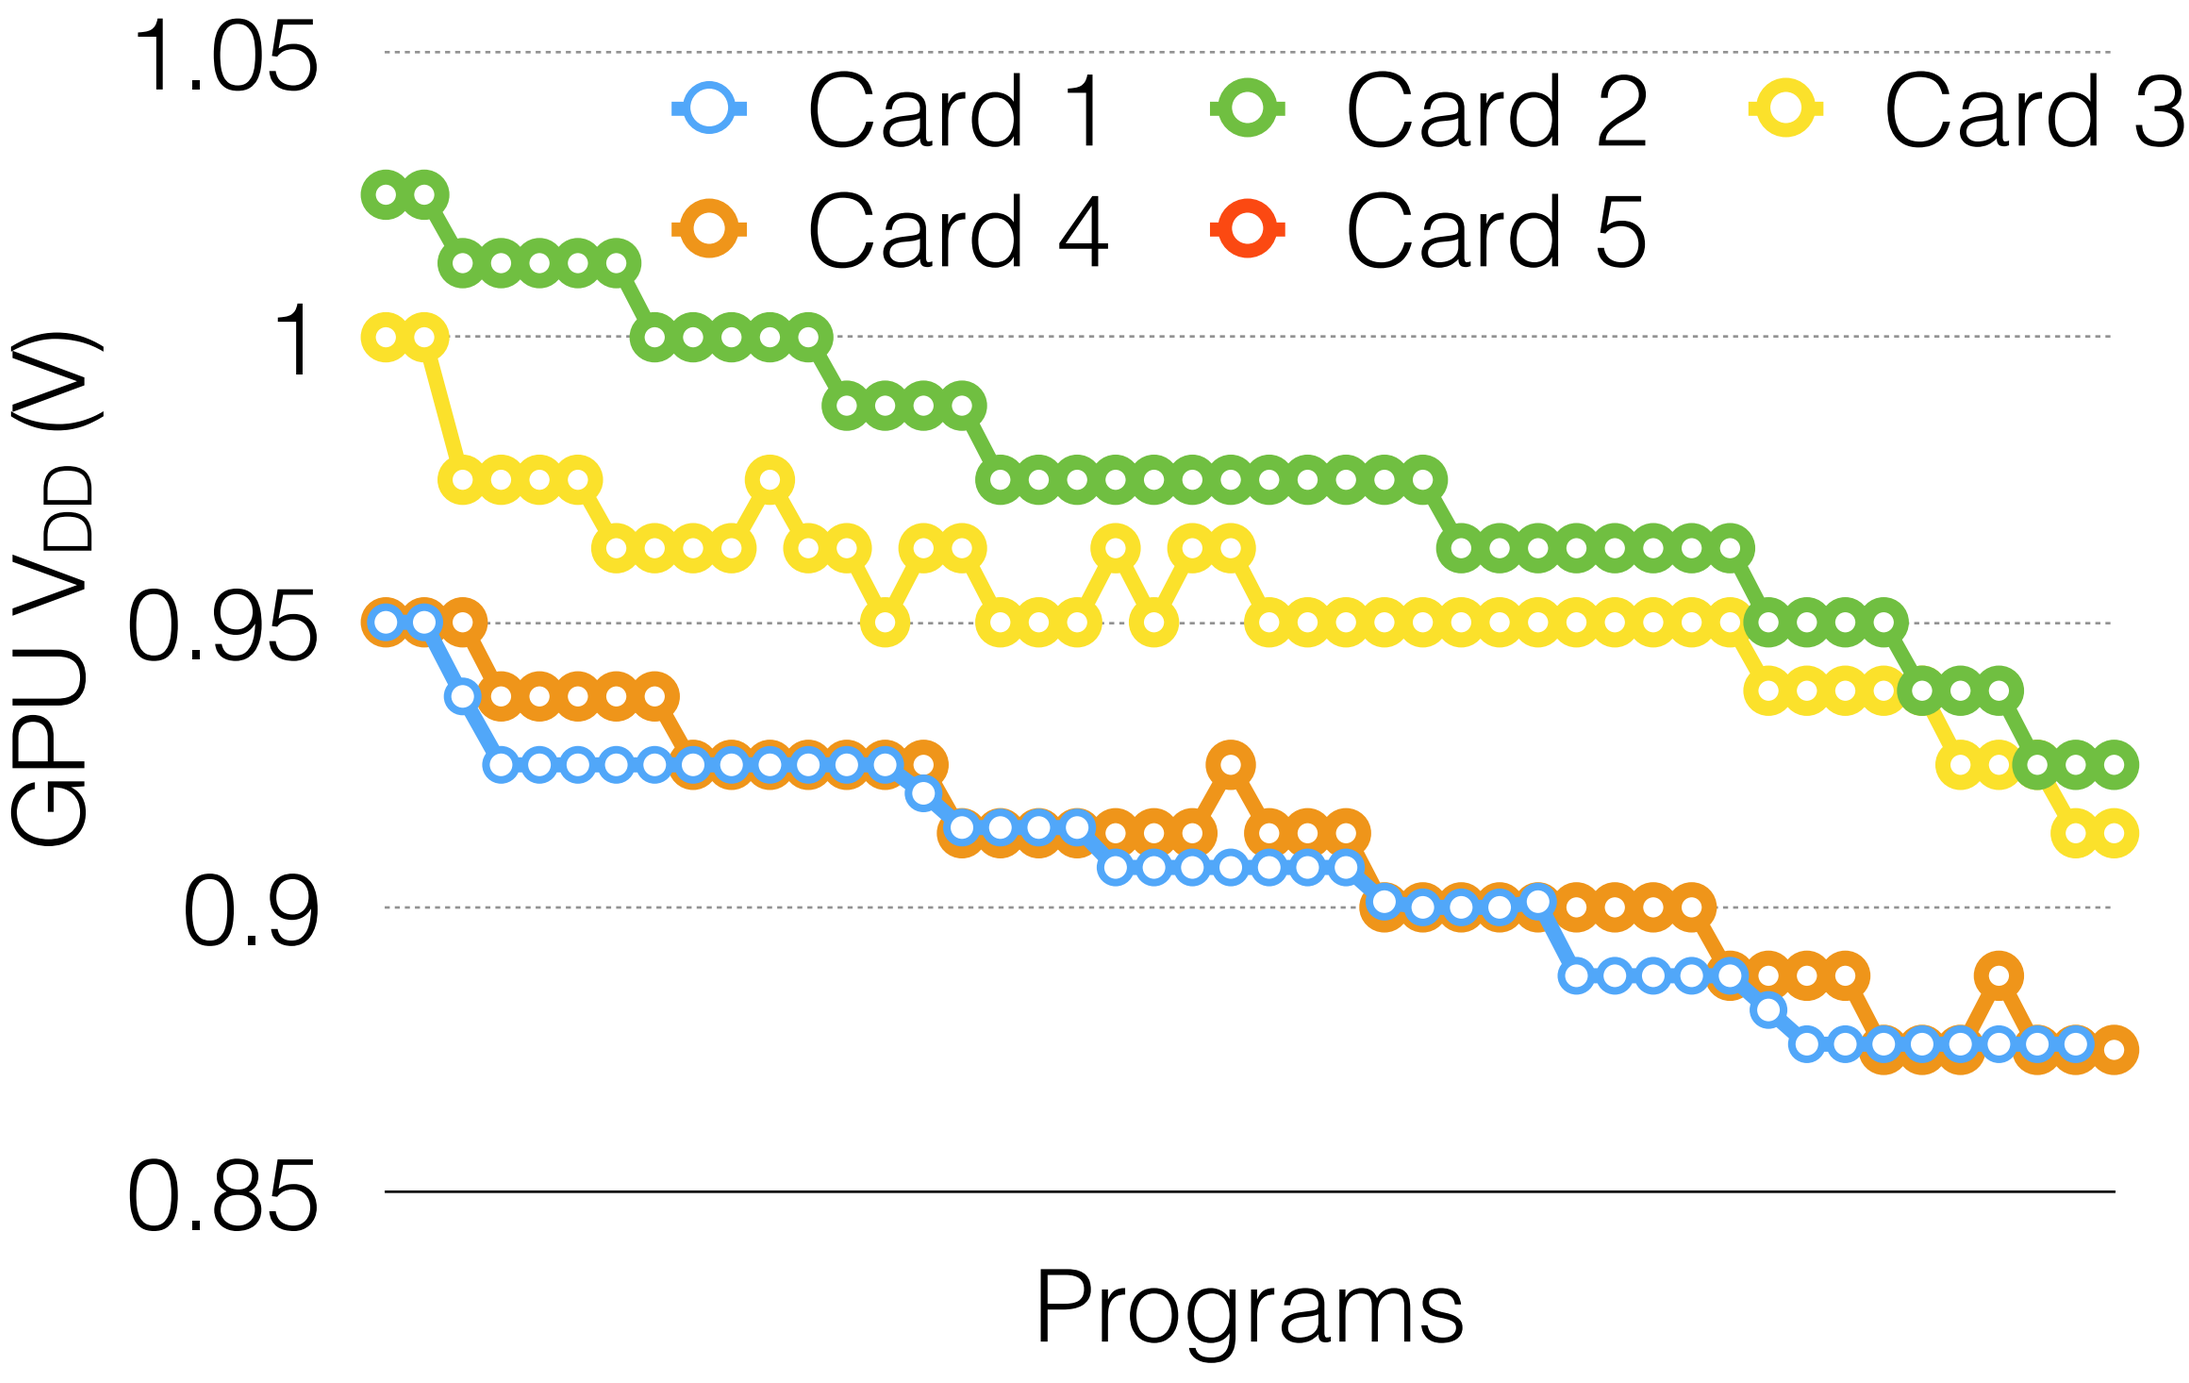

Process Variation Impact

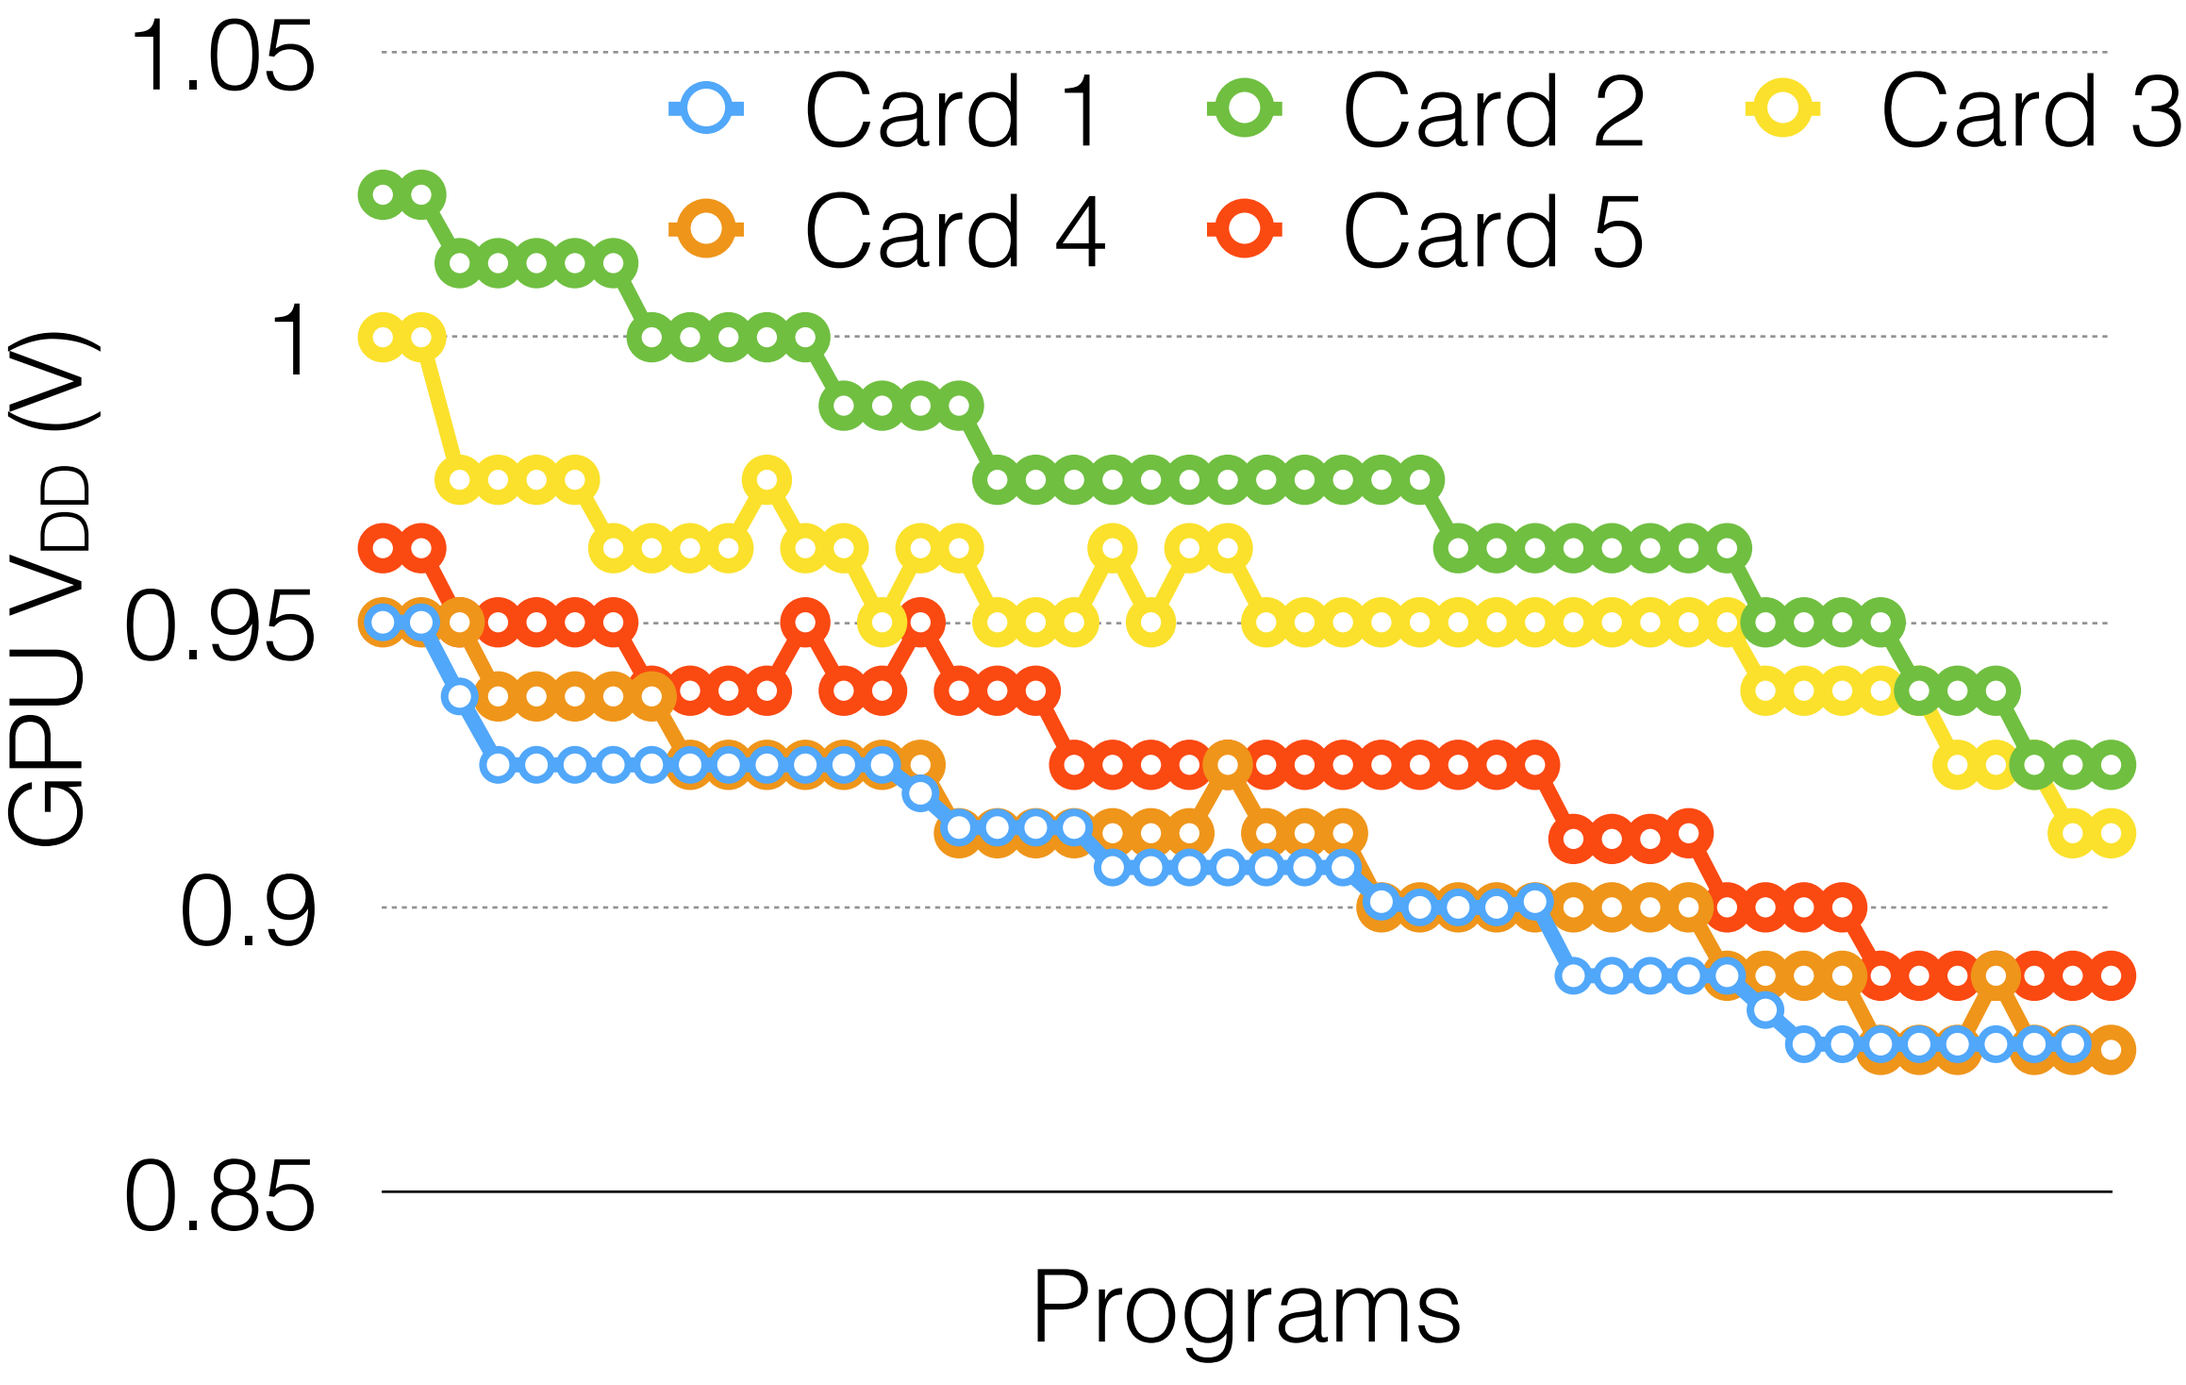

Process Variation Impact

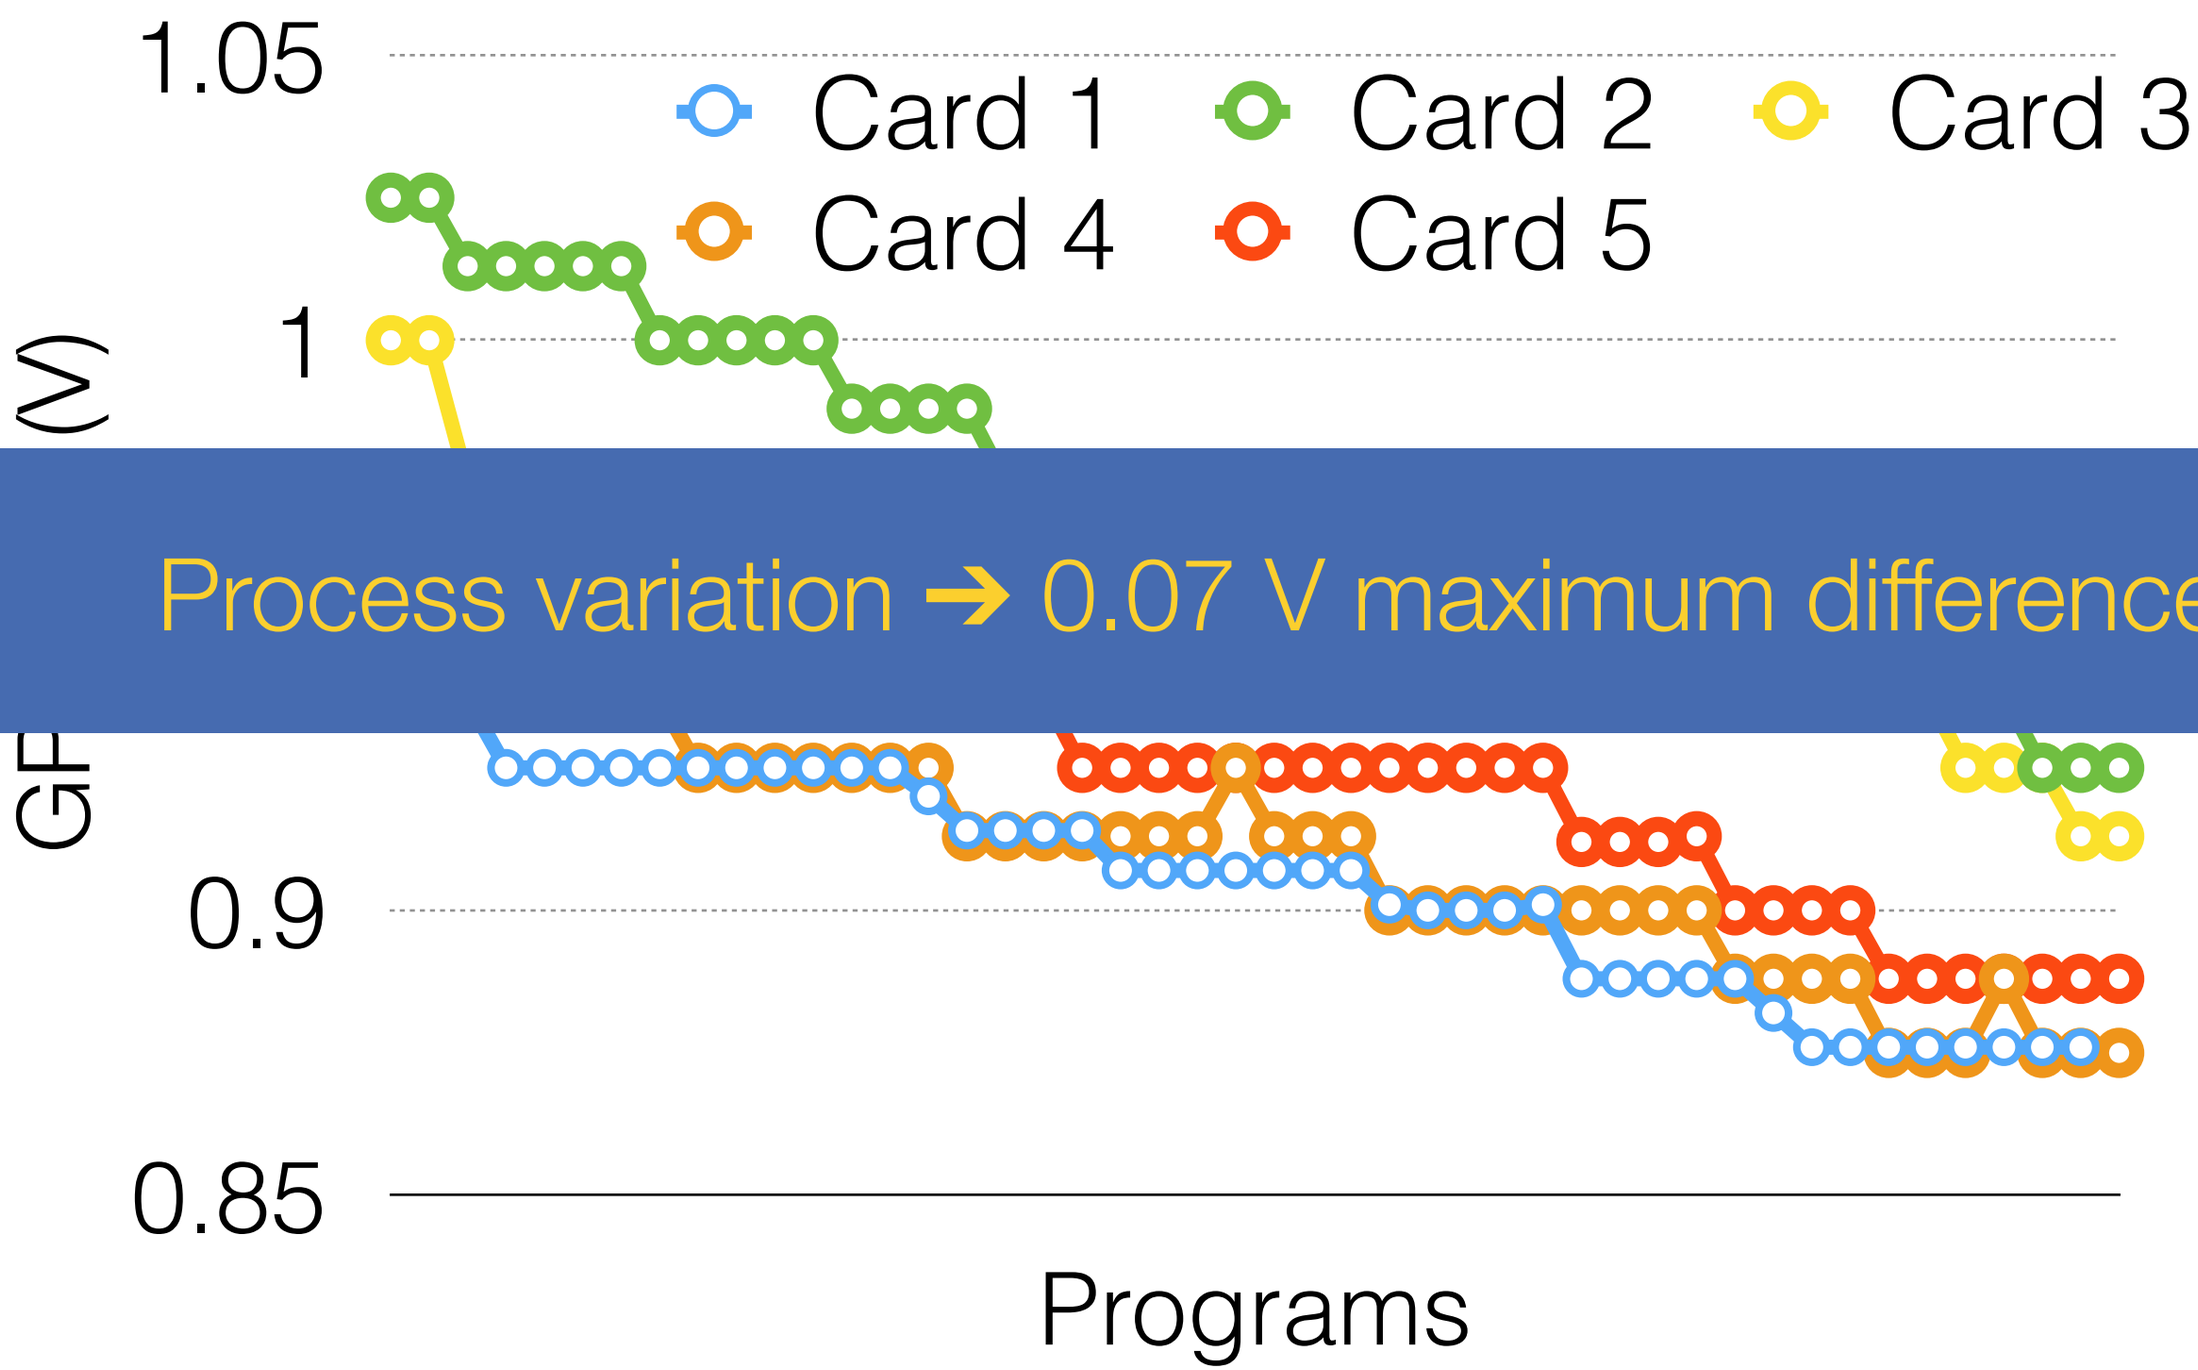

Process variation ➔ 0.07 V maximum difference

Temperature Variation Impact

Temperature Variation Impact

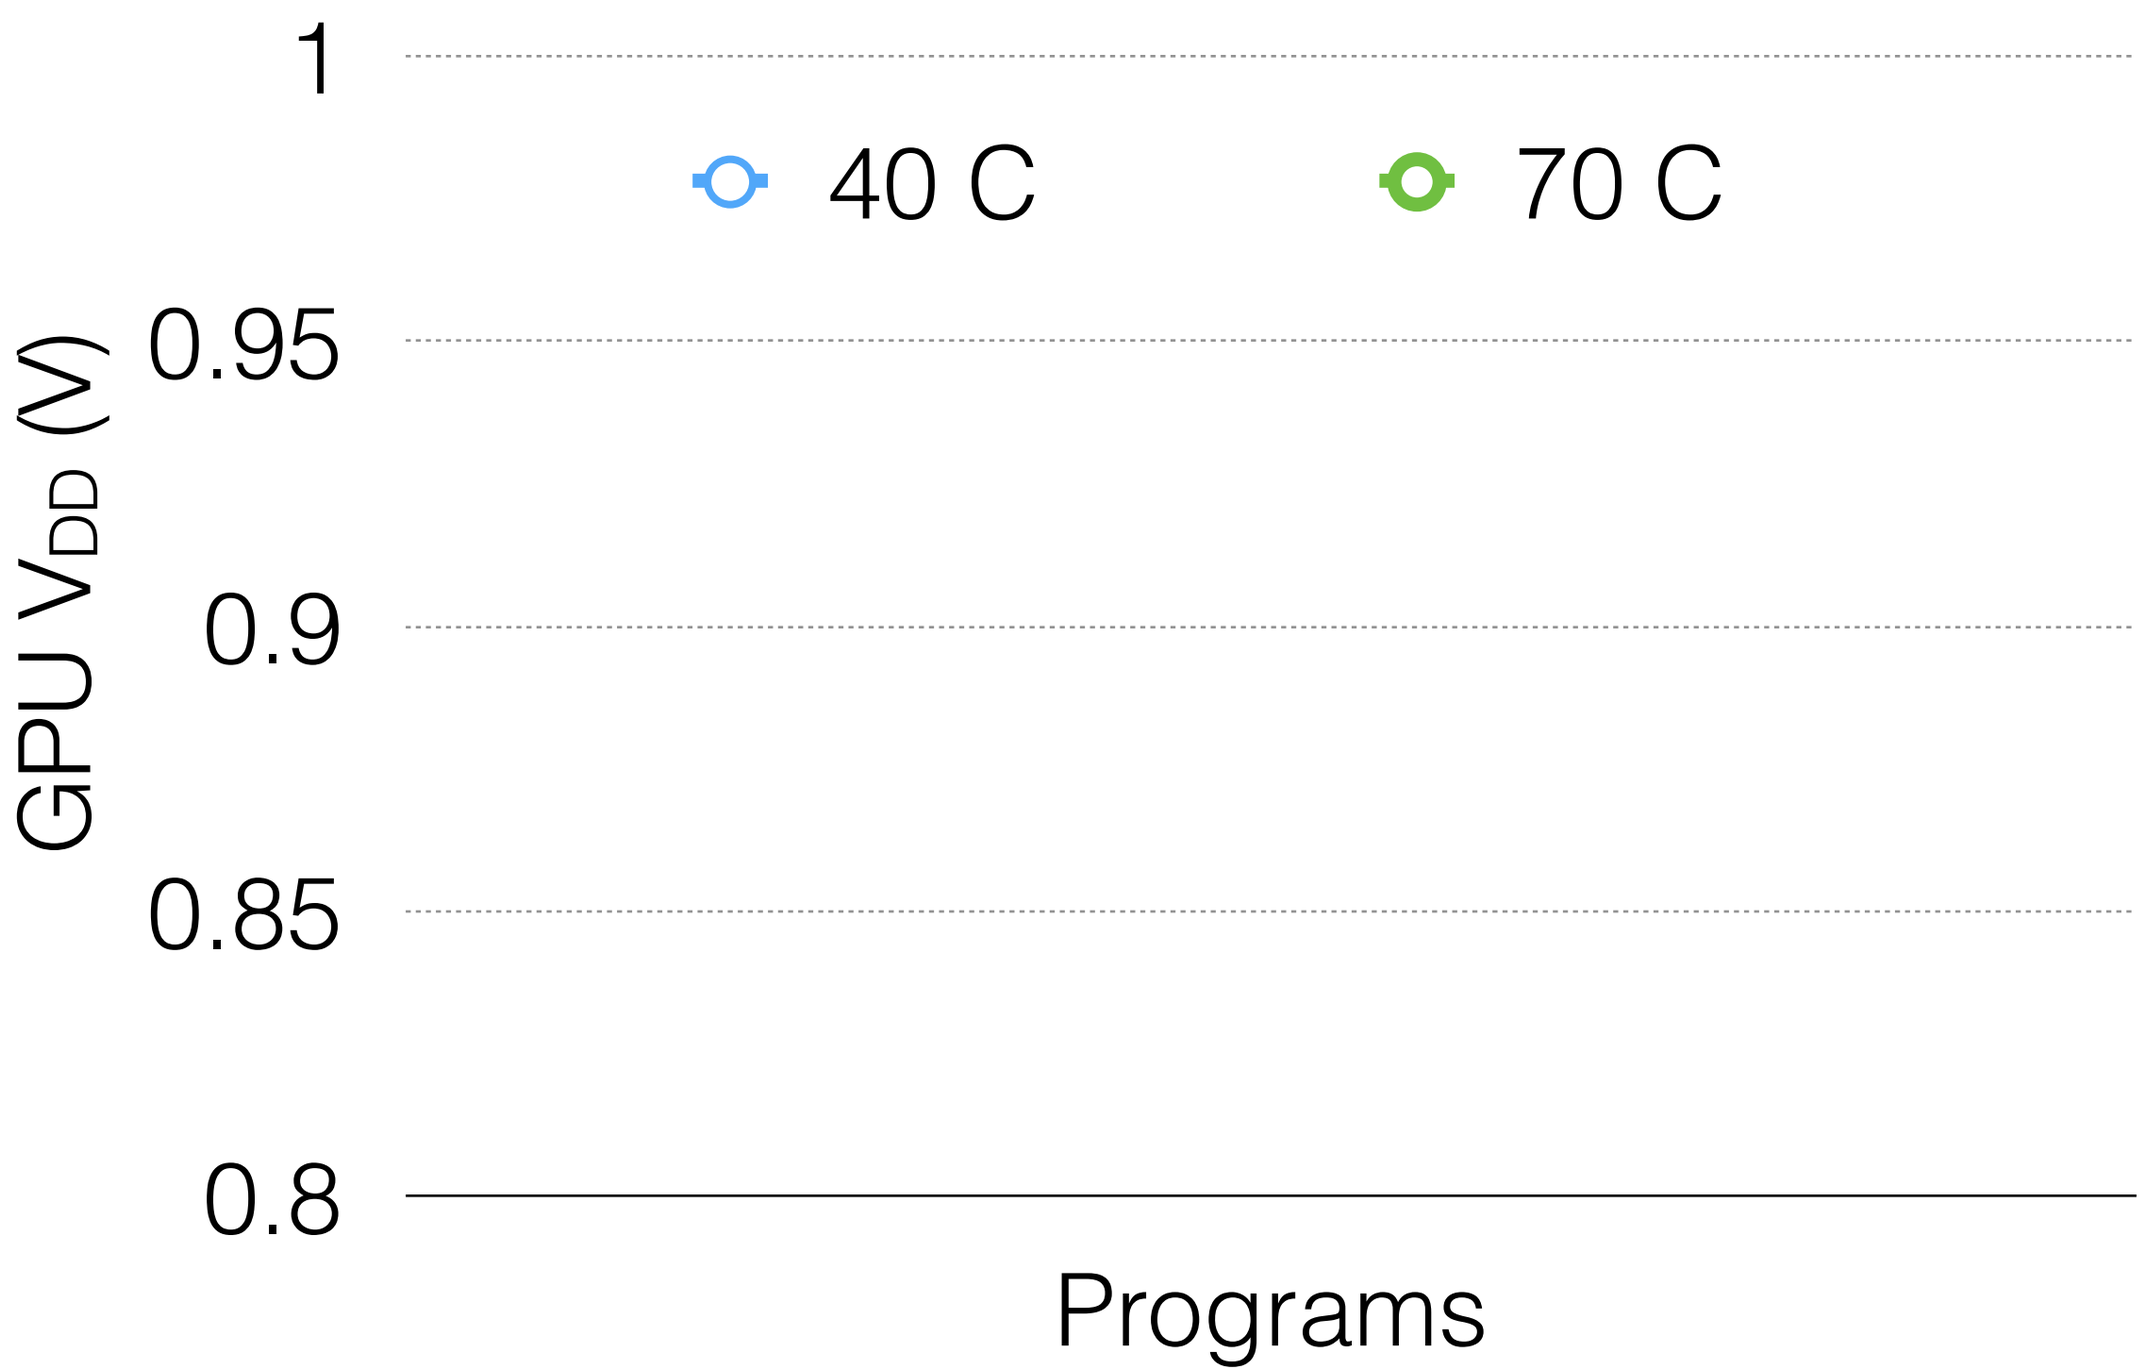

Temperature Variation Impact

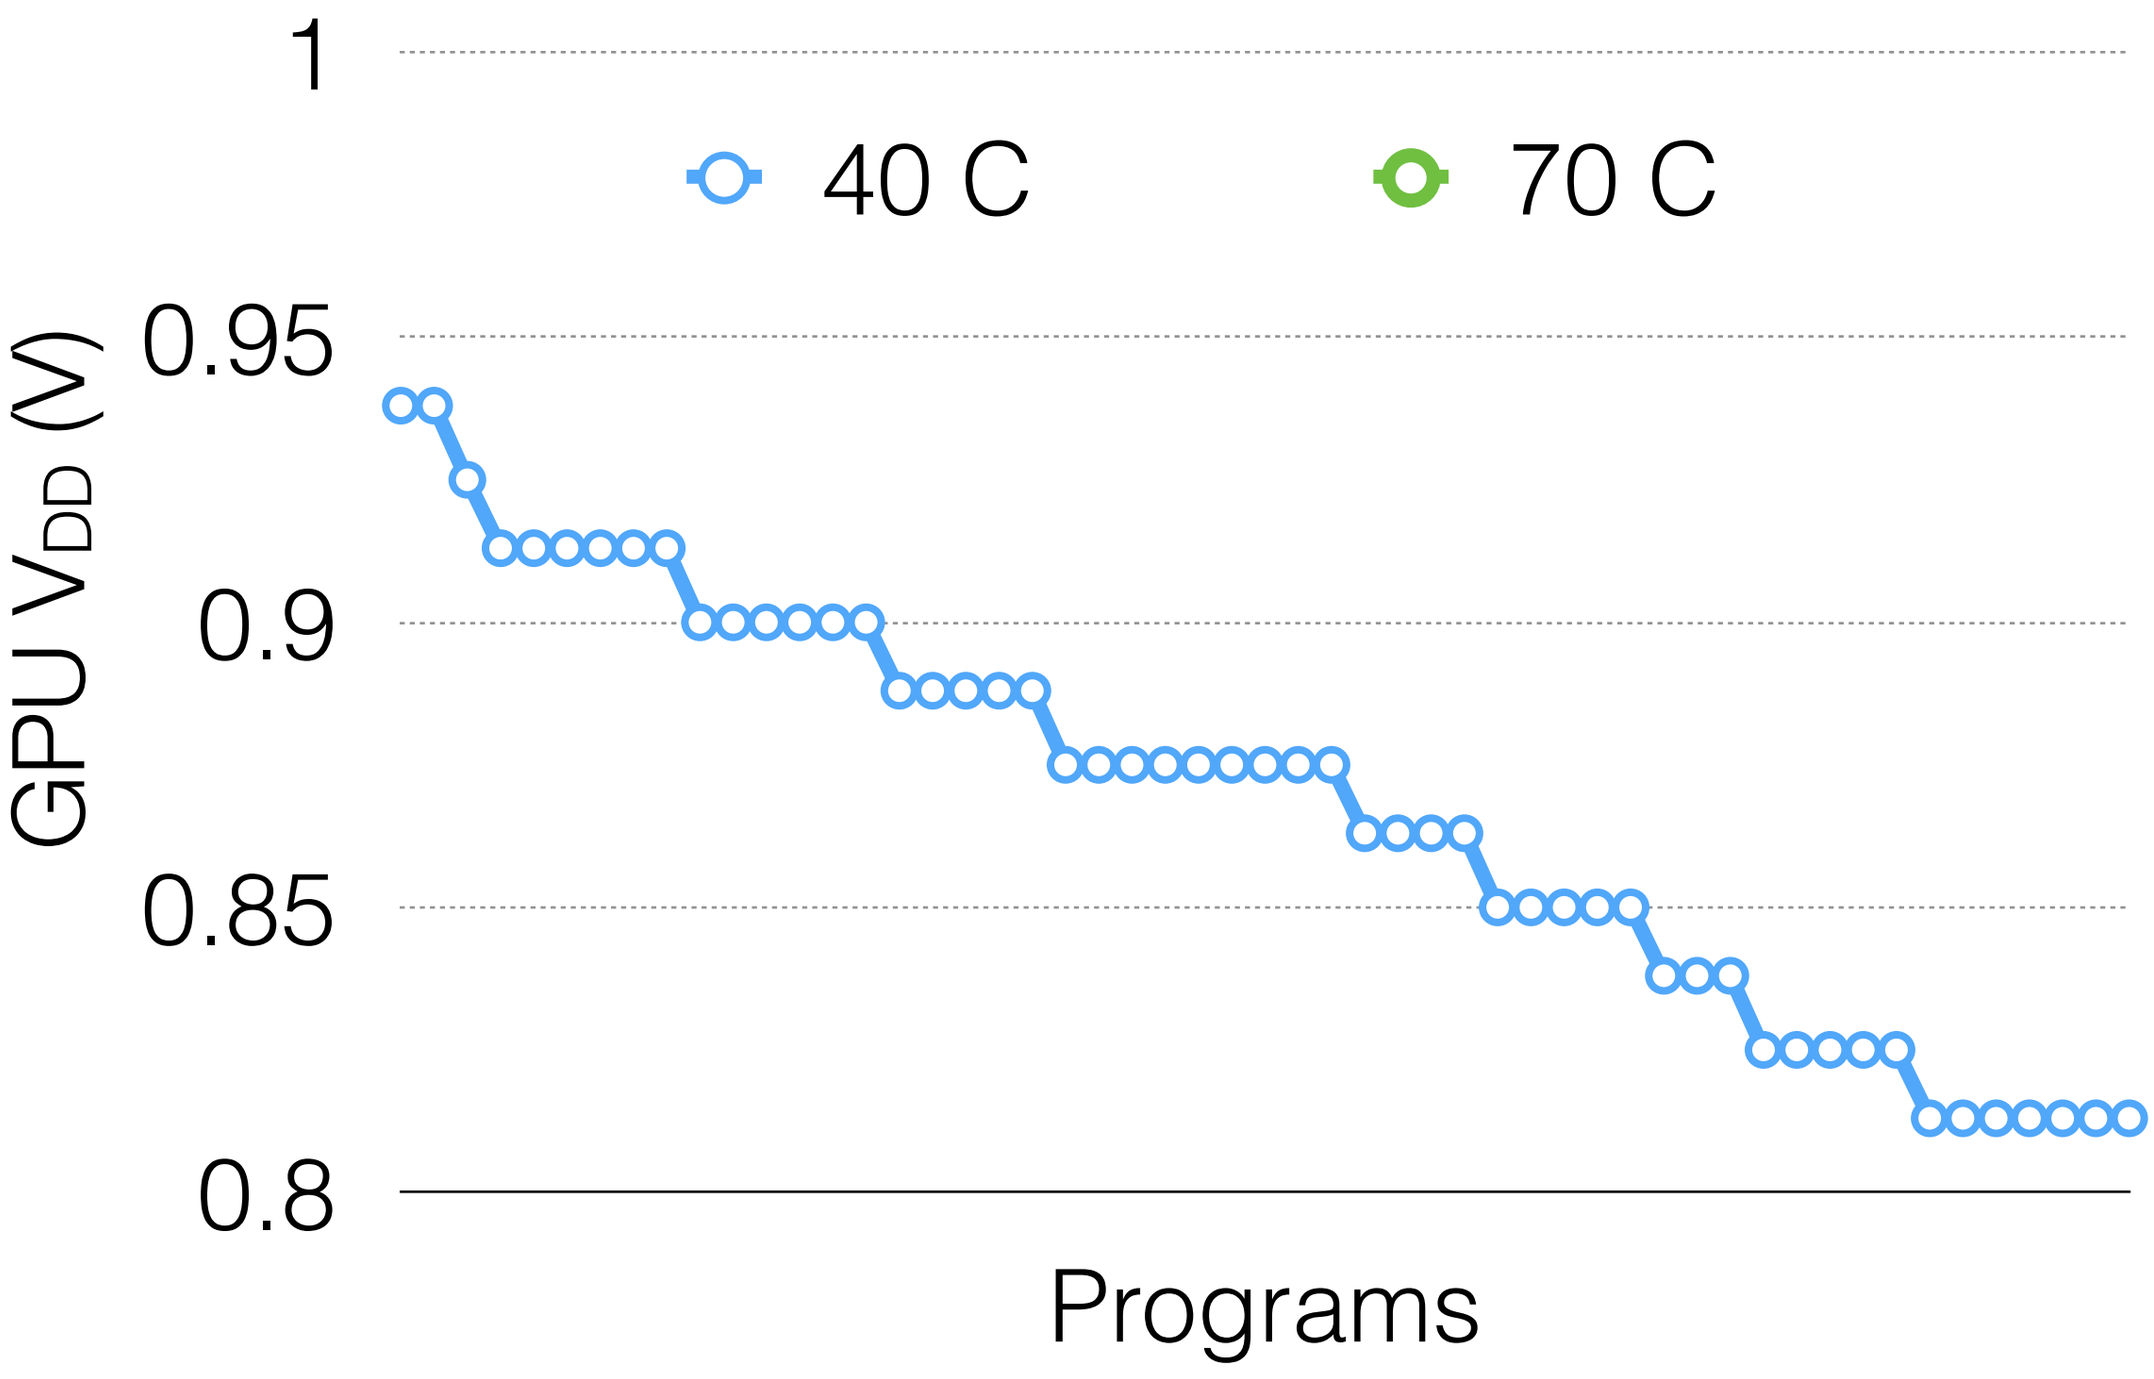

Temperature Variation Impact

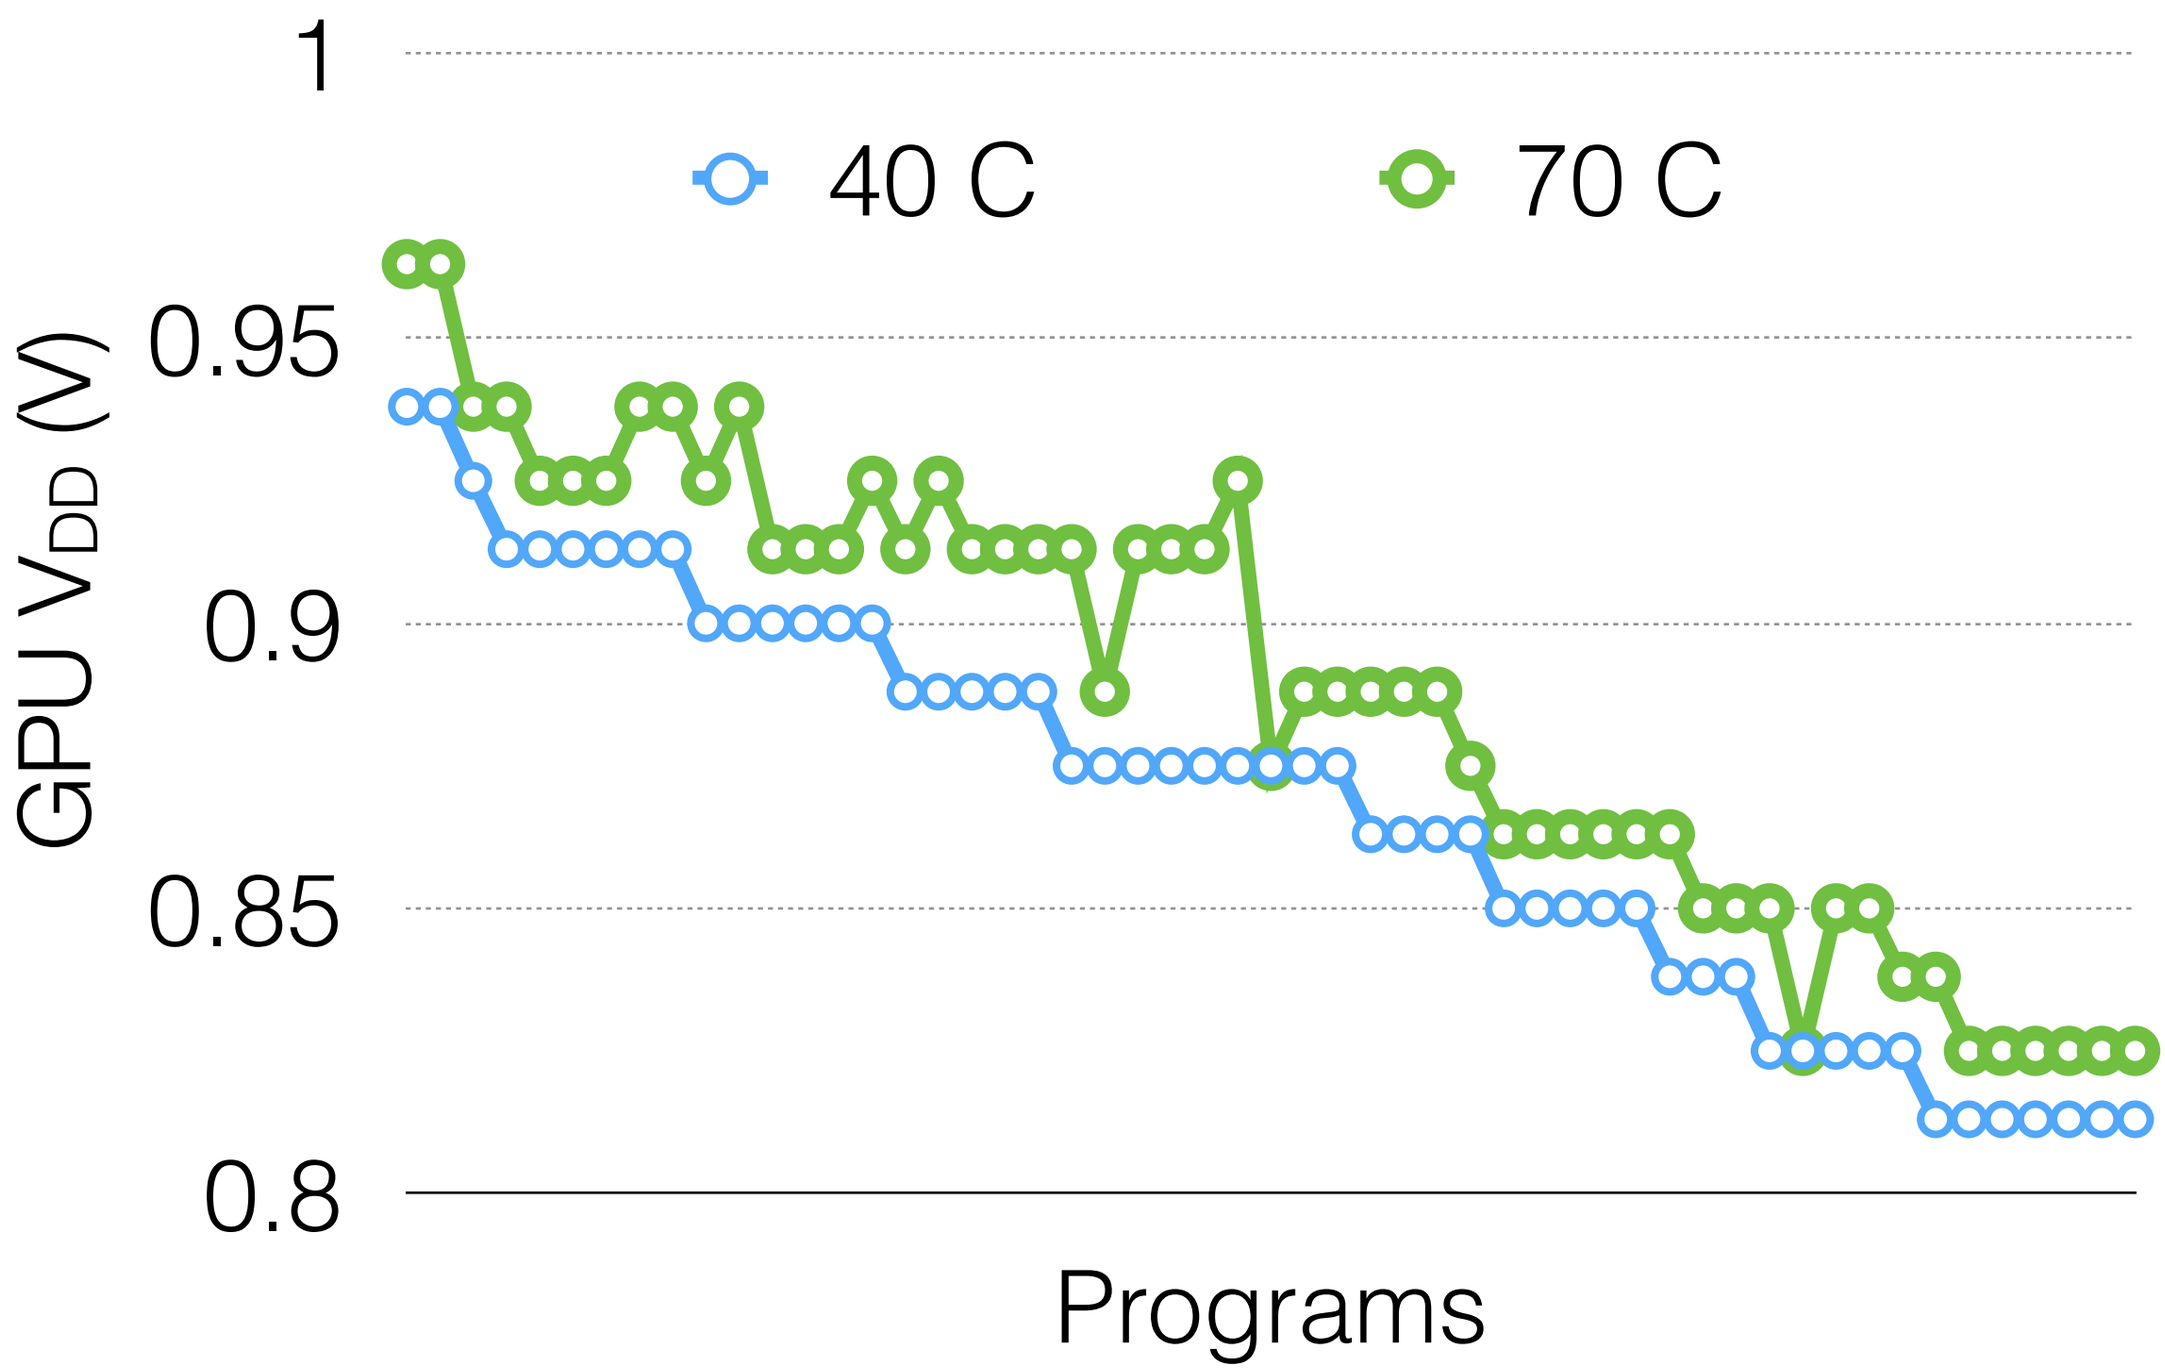

Temperature Variation Impact

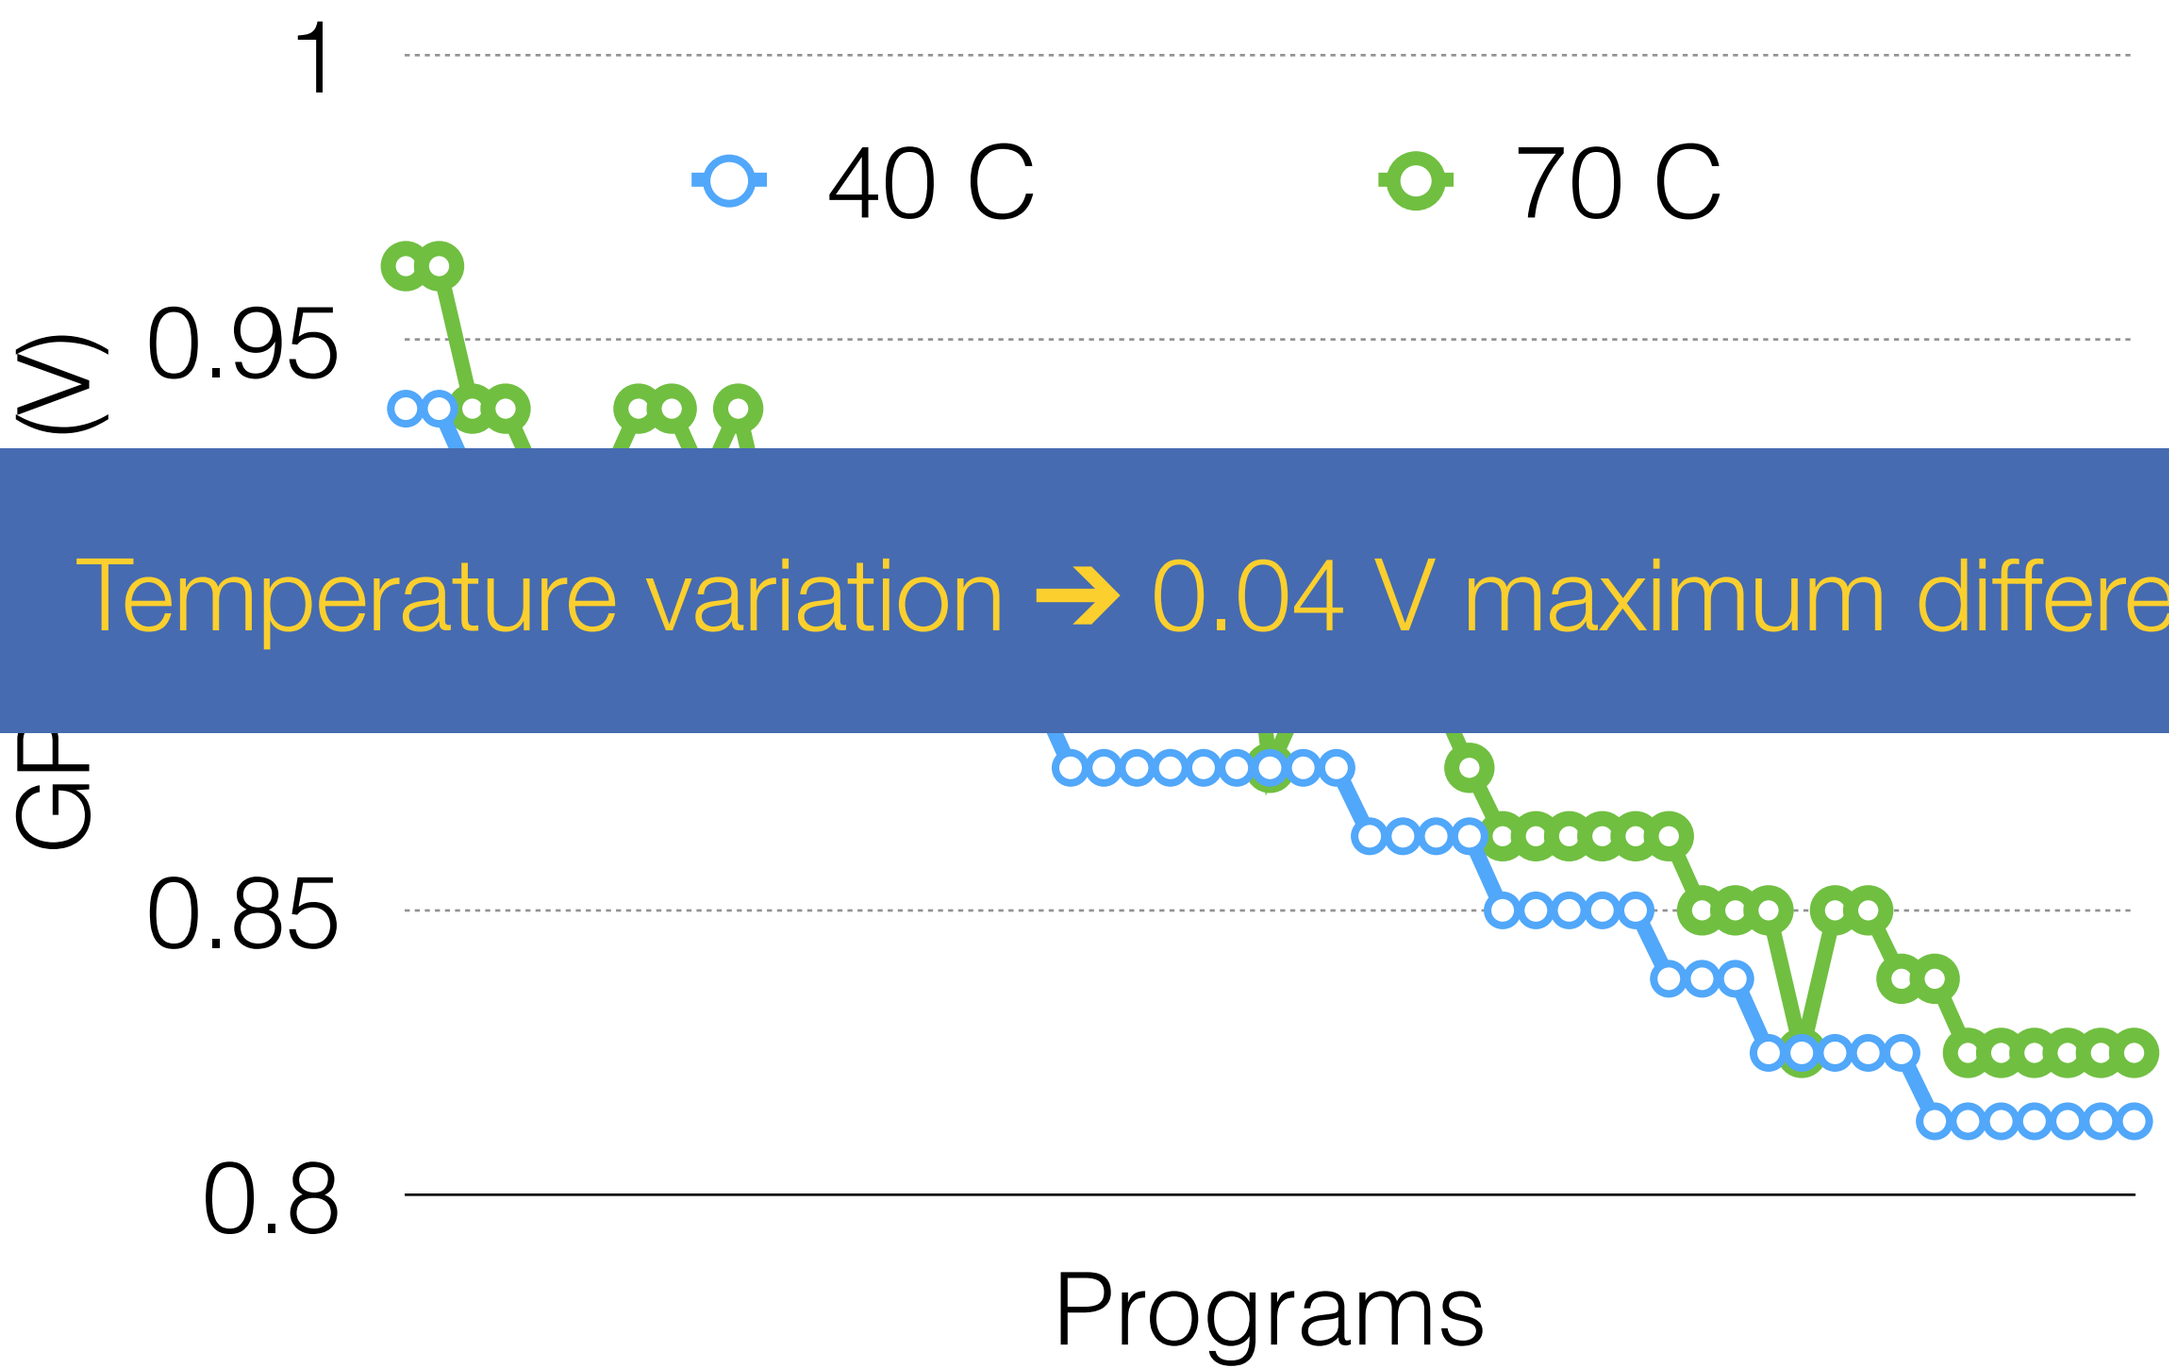

Temperature variation ➔ 0.04 V maximum difference

Combined PVT Analysis

Combined PVT Analysis

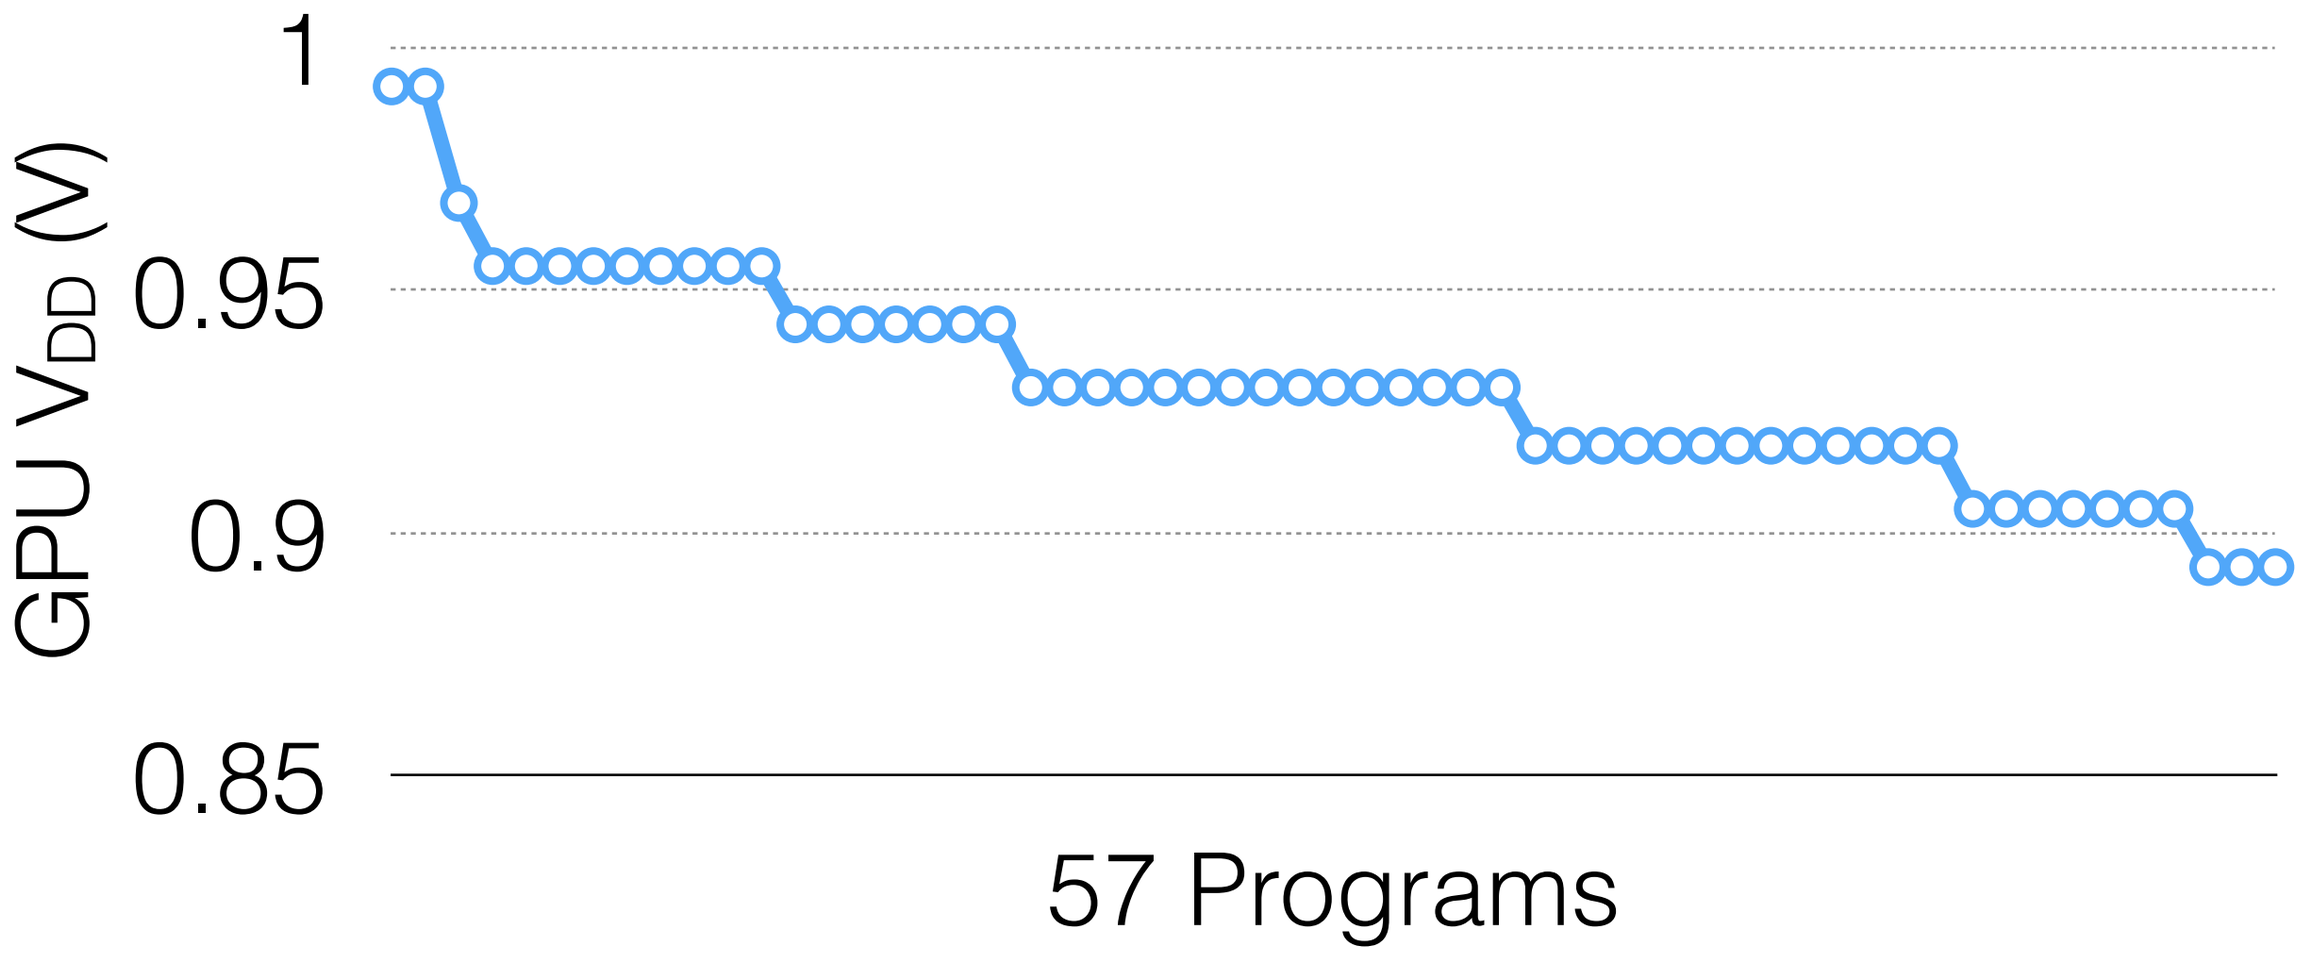

Combined PVT Analysis

- Process and temperature variation ➔ relatively uniform impact on ALL programs

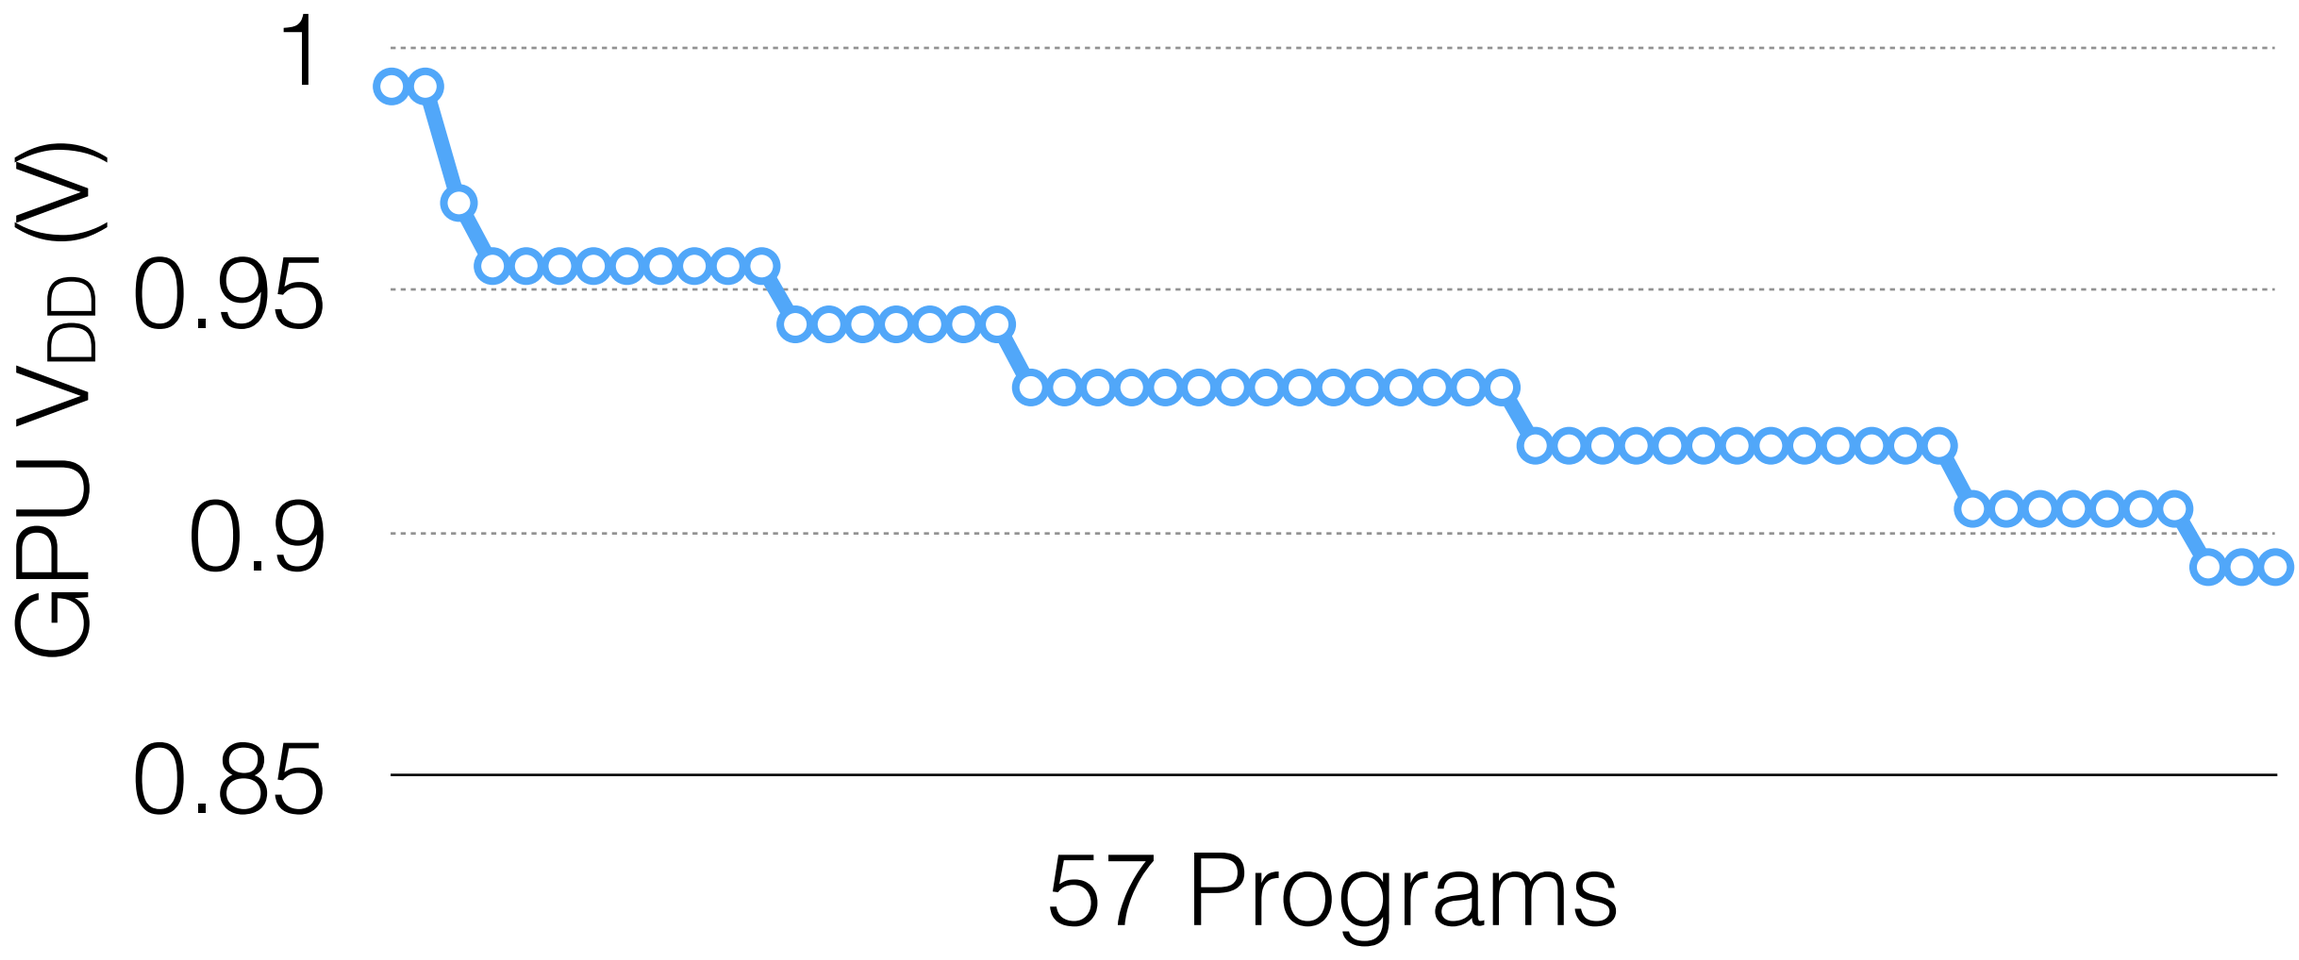

Combined PVT Analysis

- Process and temperature variation ➔ relatively uniform impact on ALL programs

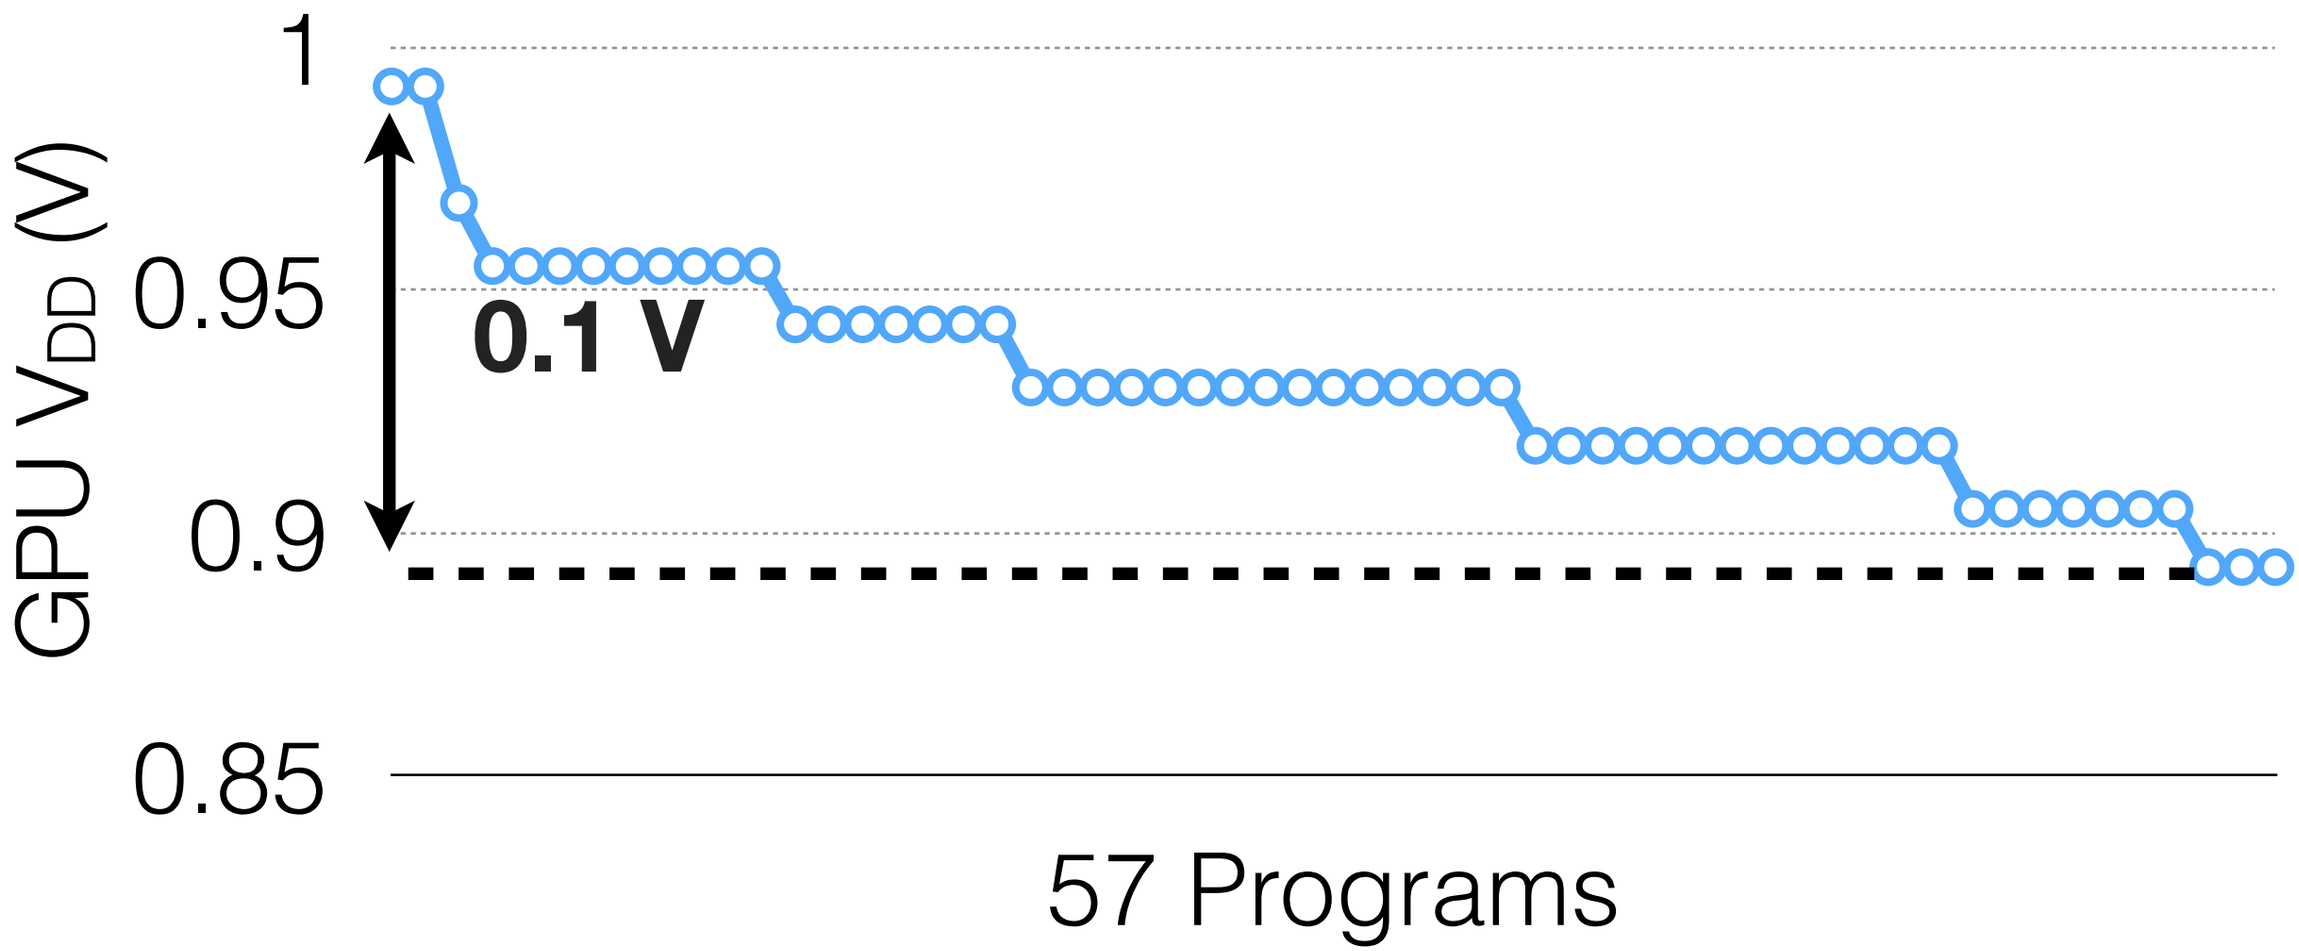

Combined PVT Analysis

- Process and temperature variation ➔ relatively uniform impact on ALL programs
- Voltage variation ➔ 0.1 V difference across programs

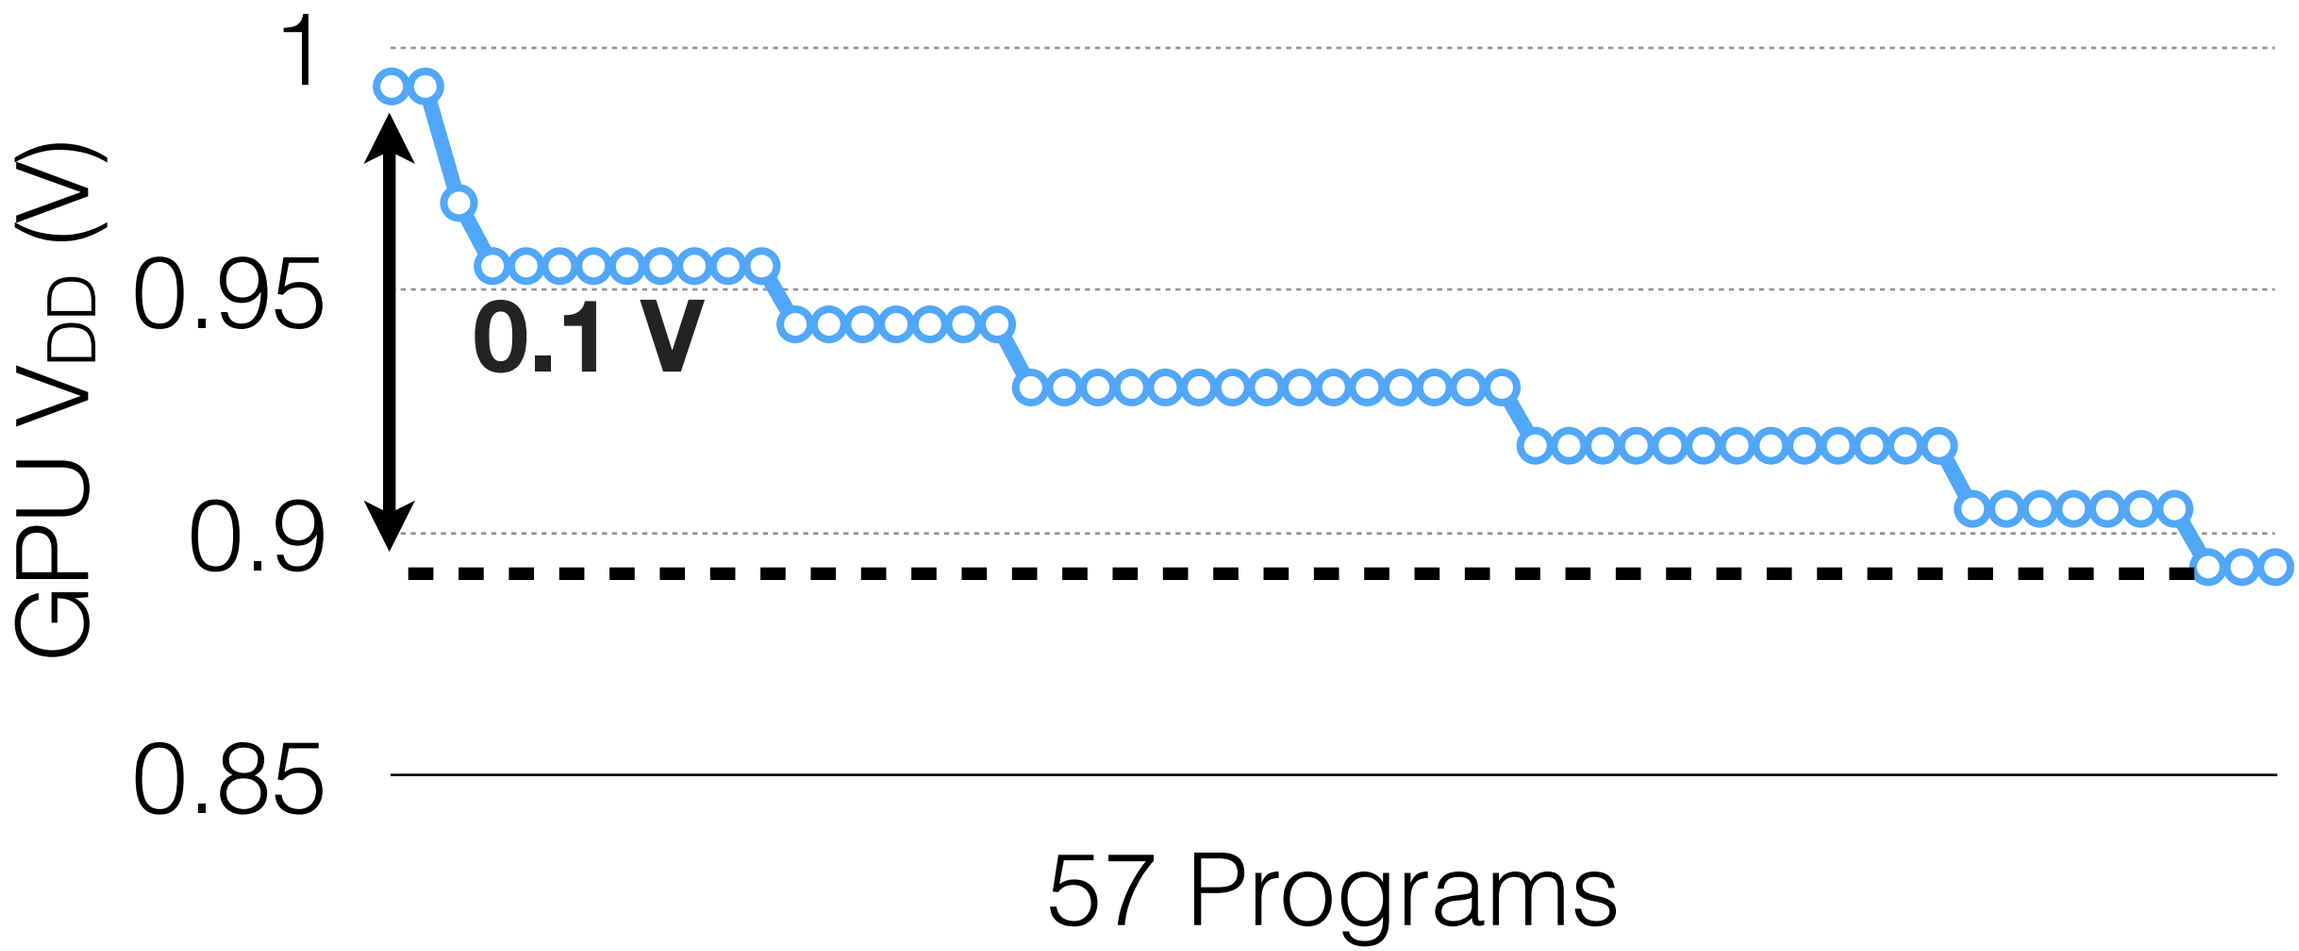

Combined PVT Analysis

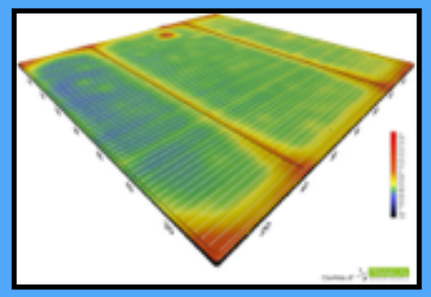

Process

?%

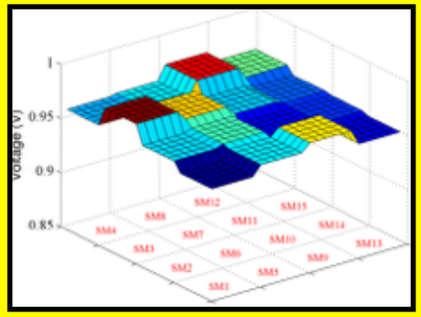

Voltage

?%

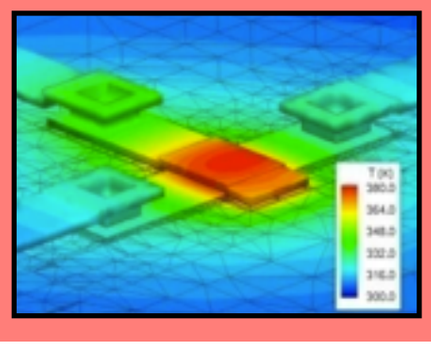

Tempera
-ture

?%

Combined PVT Analysis

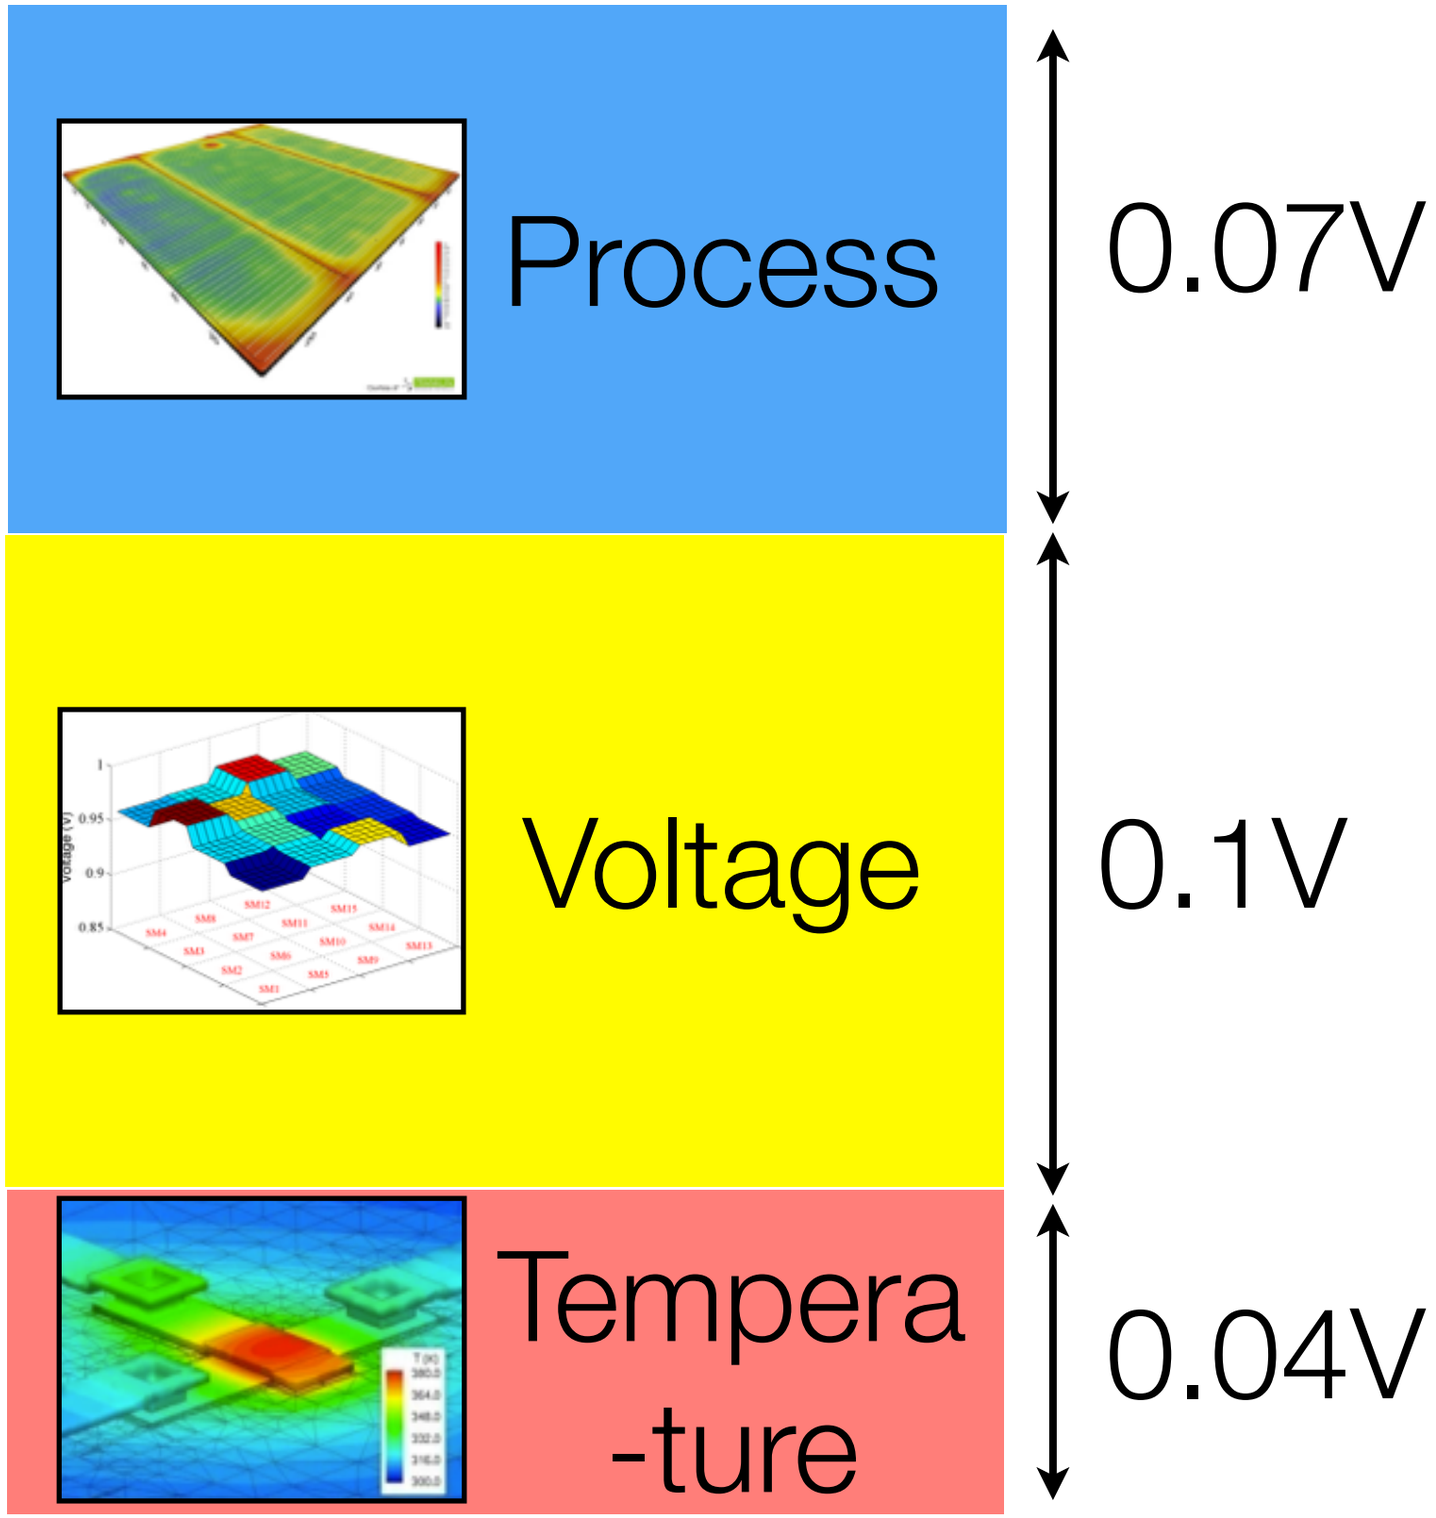

Voltage Noise Background

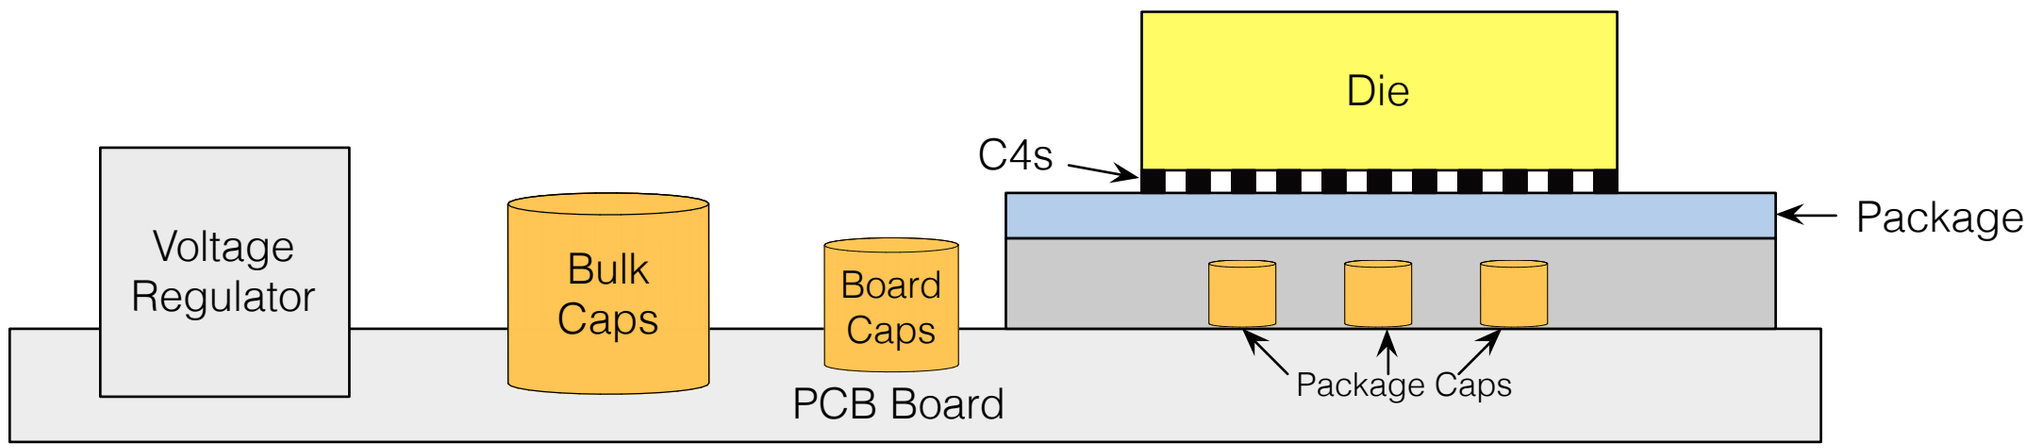

Voltage Noise Background

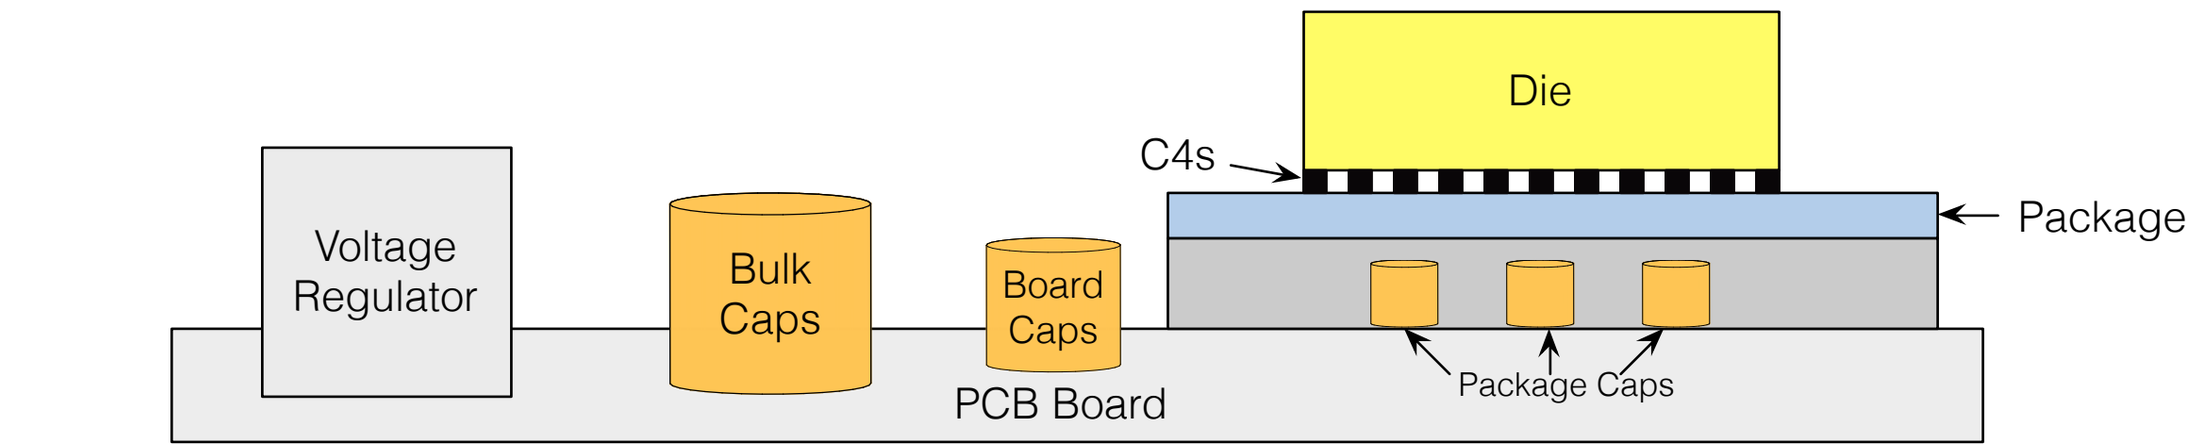

Voltage Noise Background

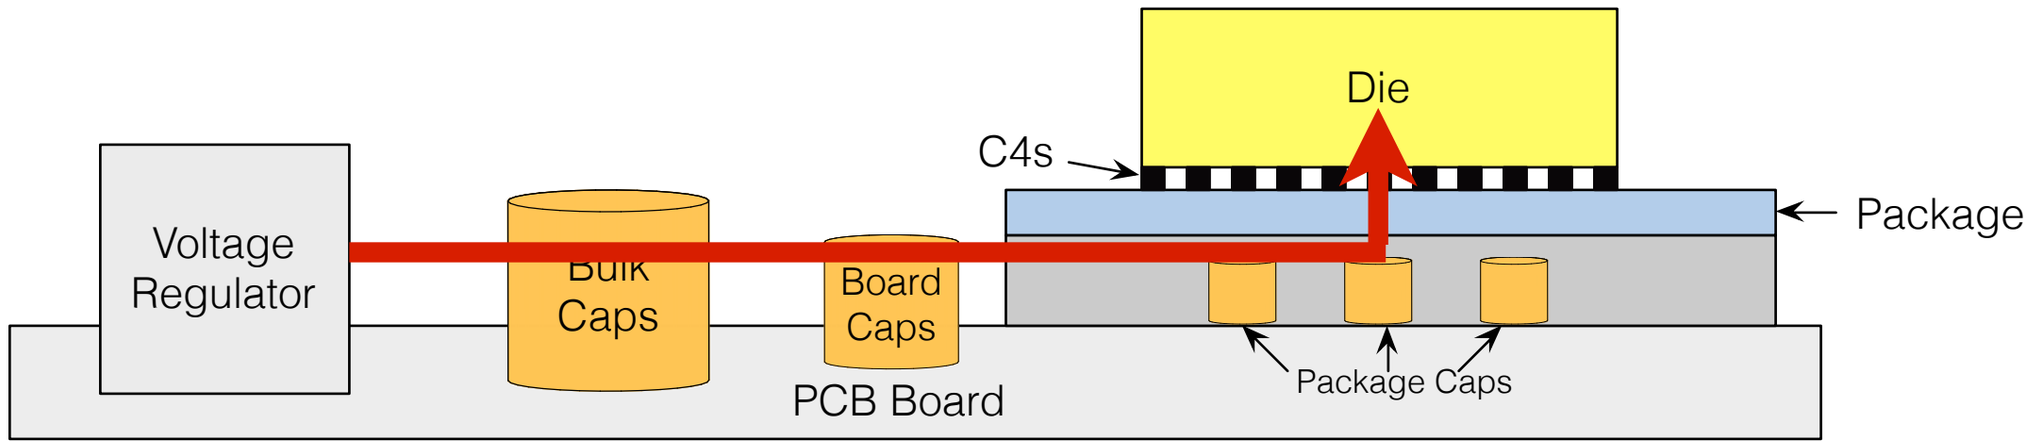

Voltage Noise Background

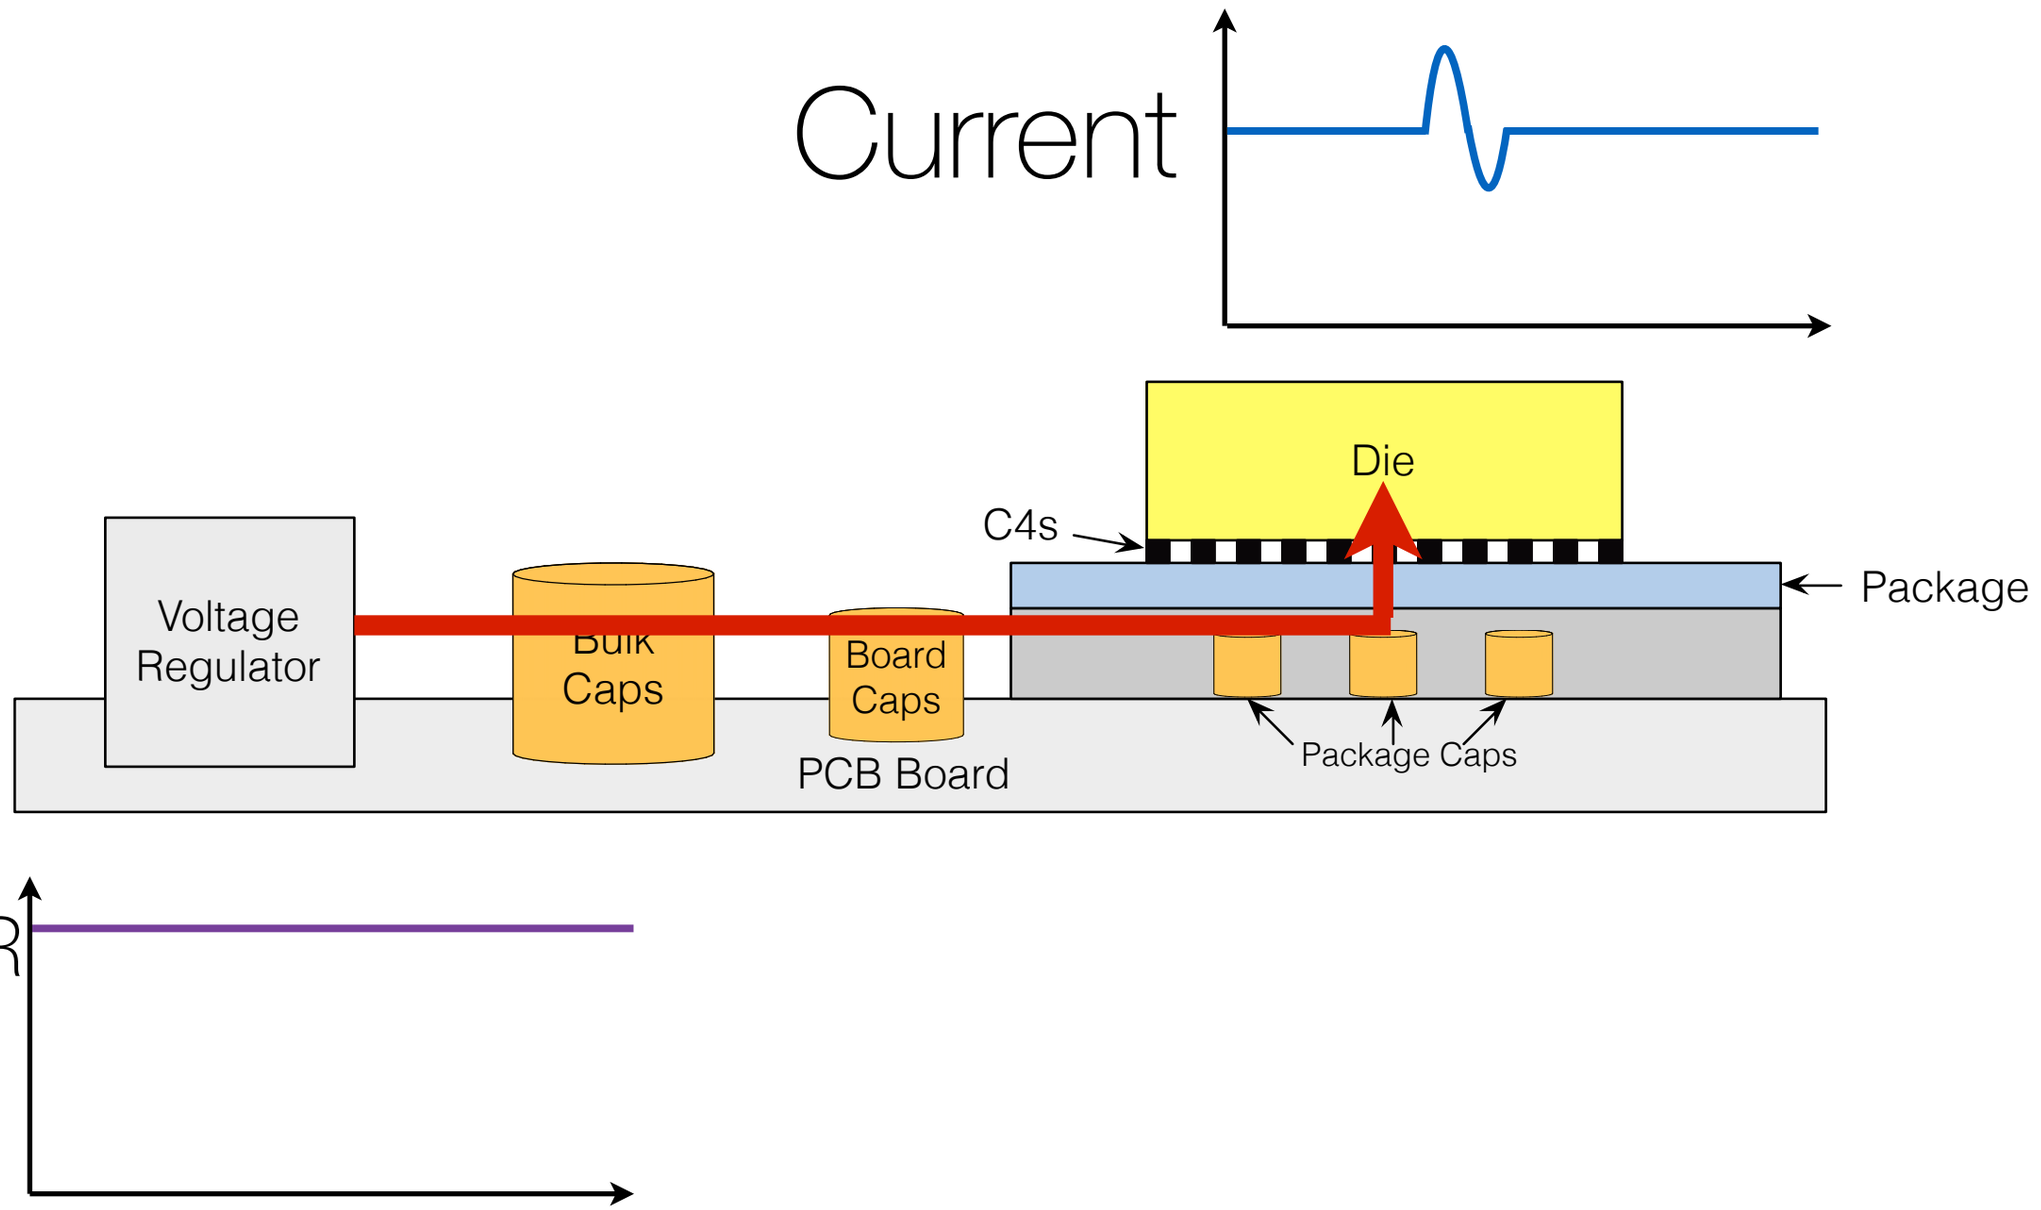

Voltage Noise Background

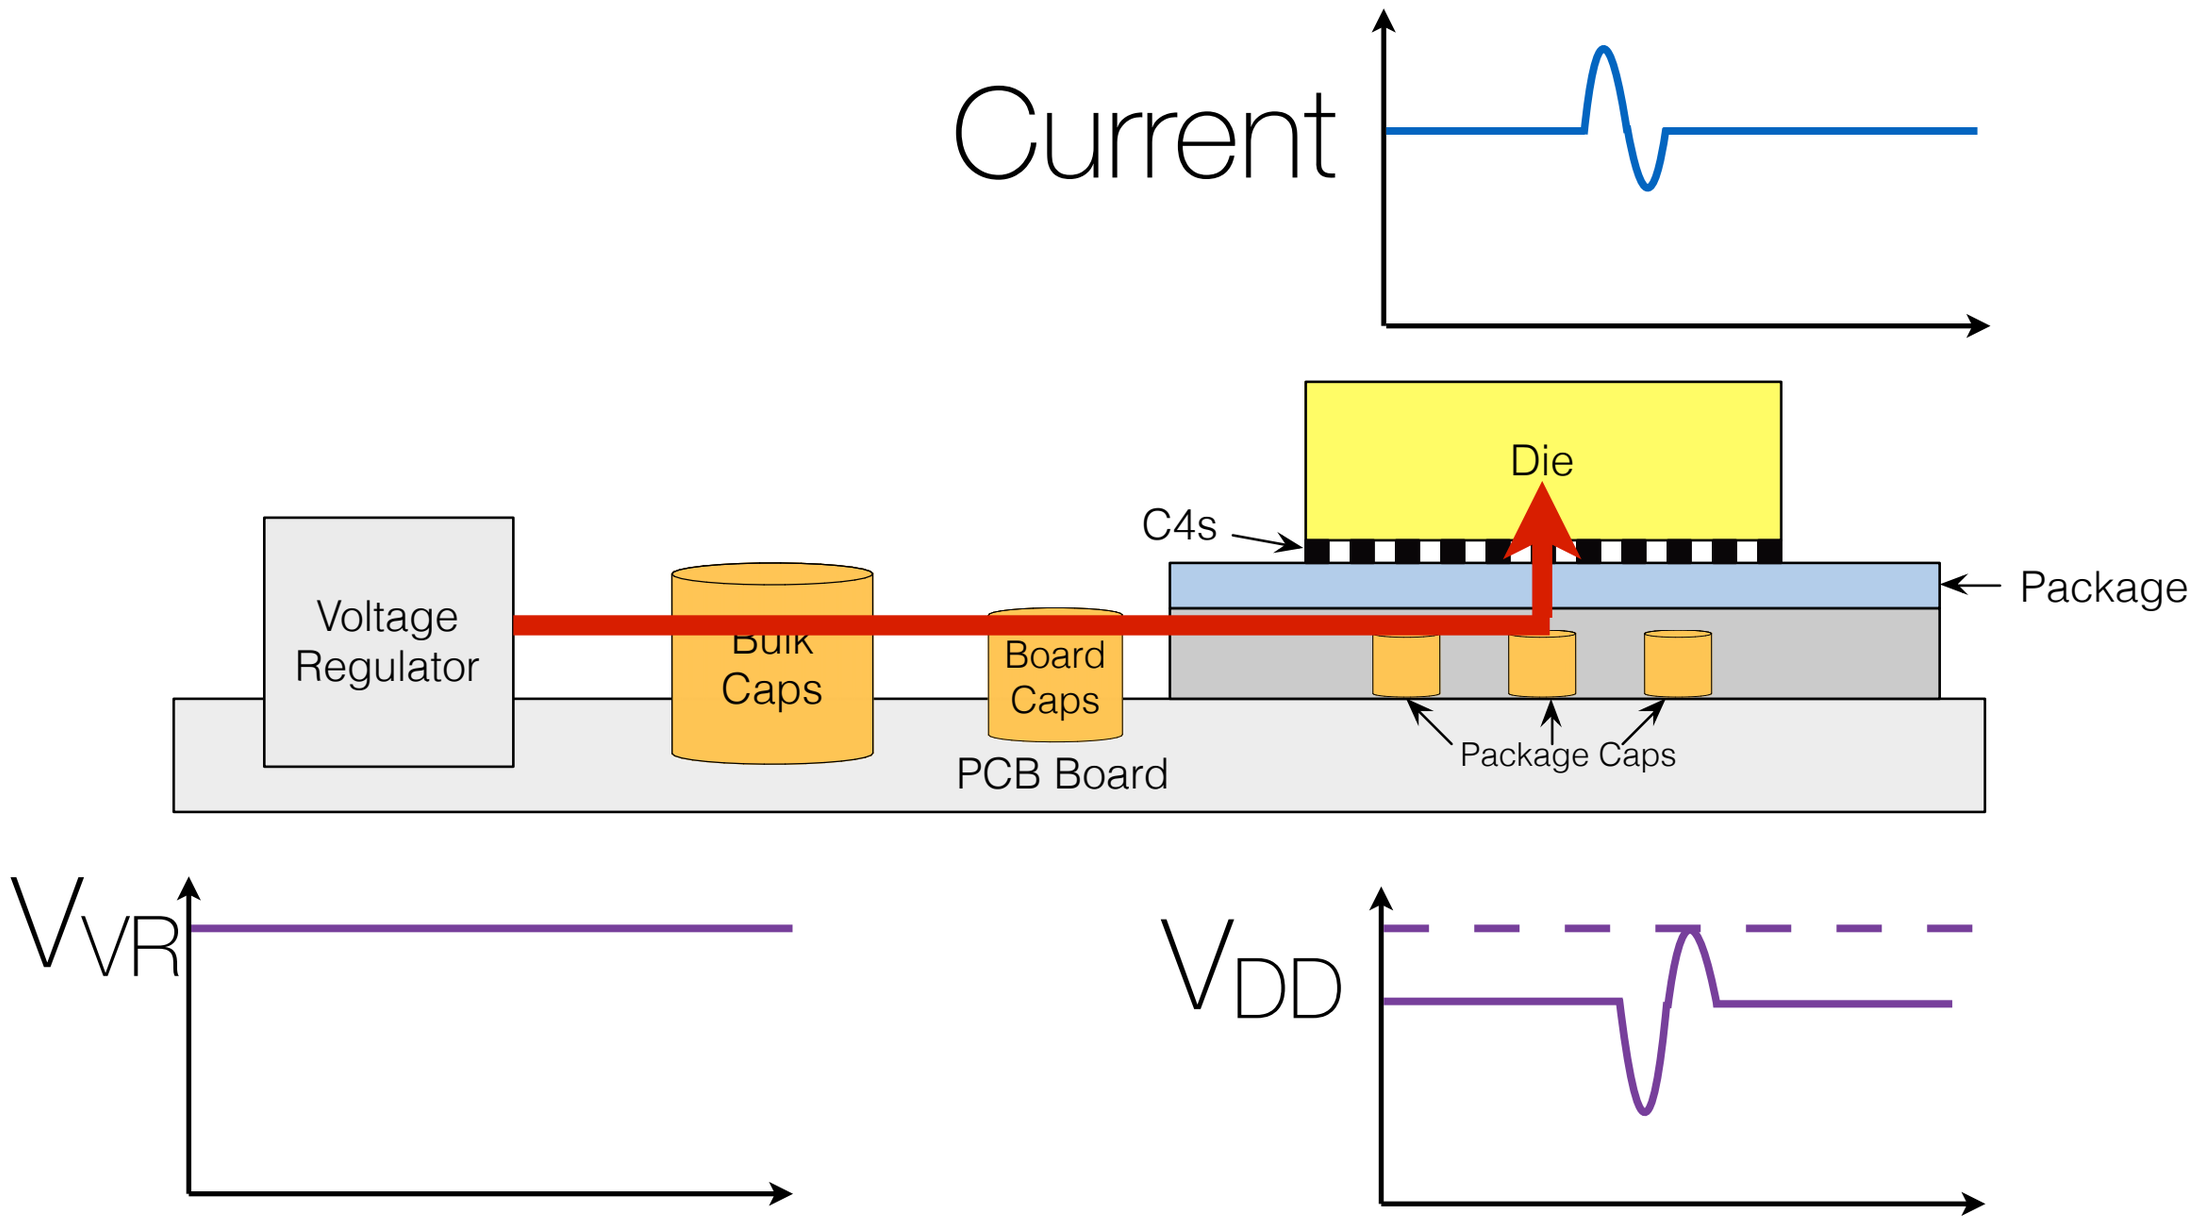

Voltage Noise Background

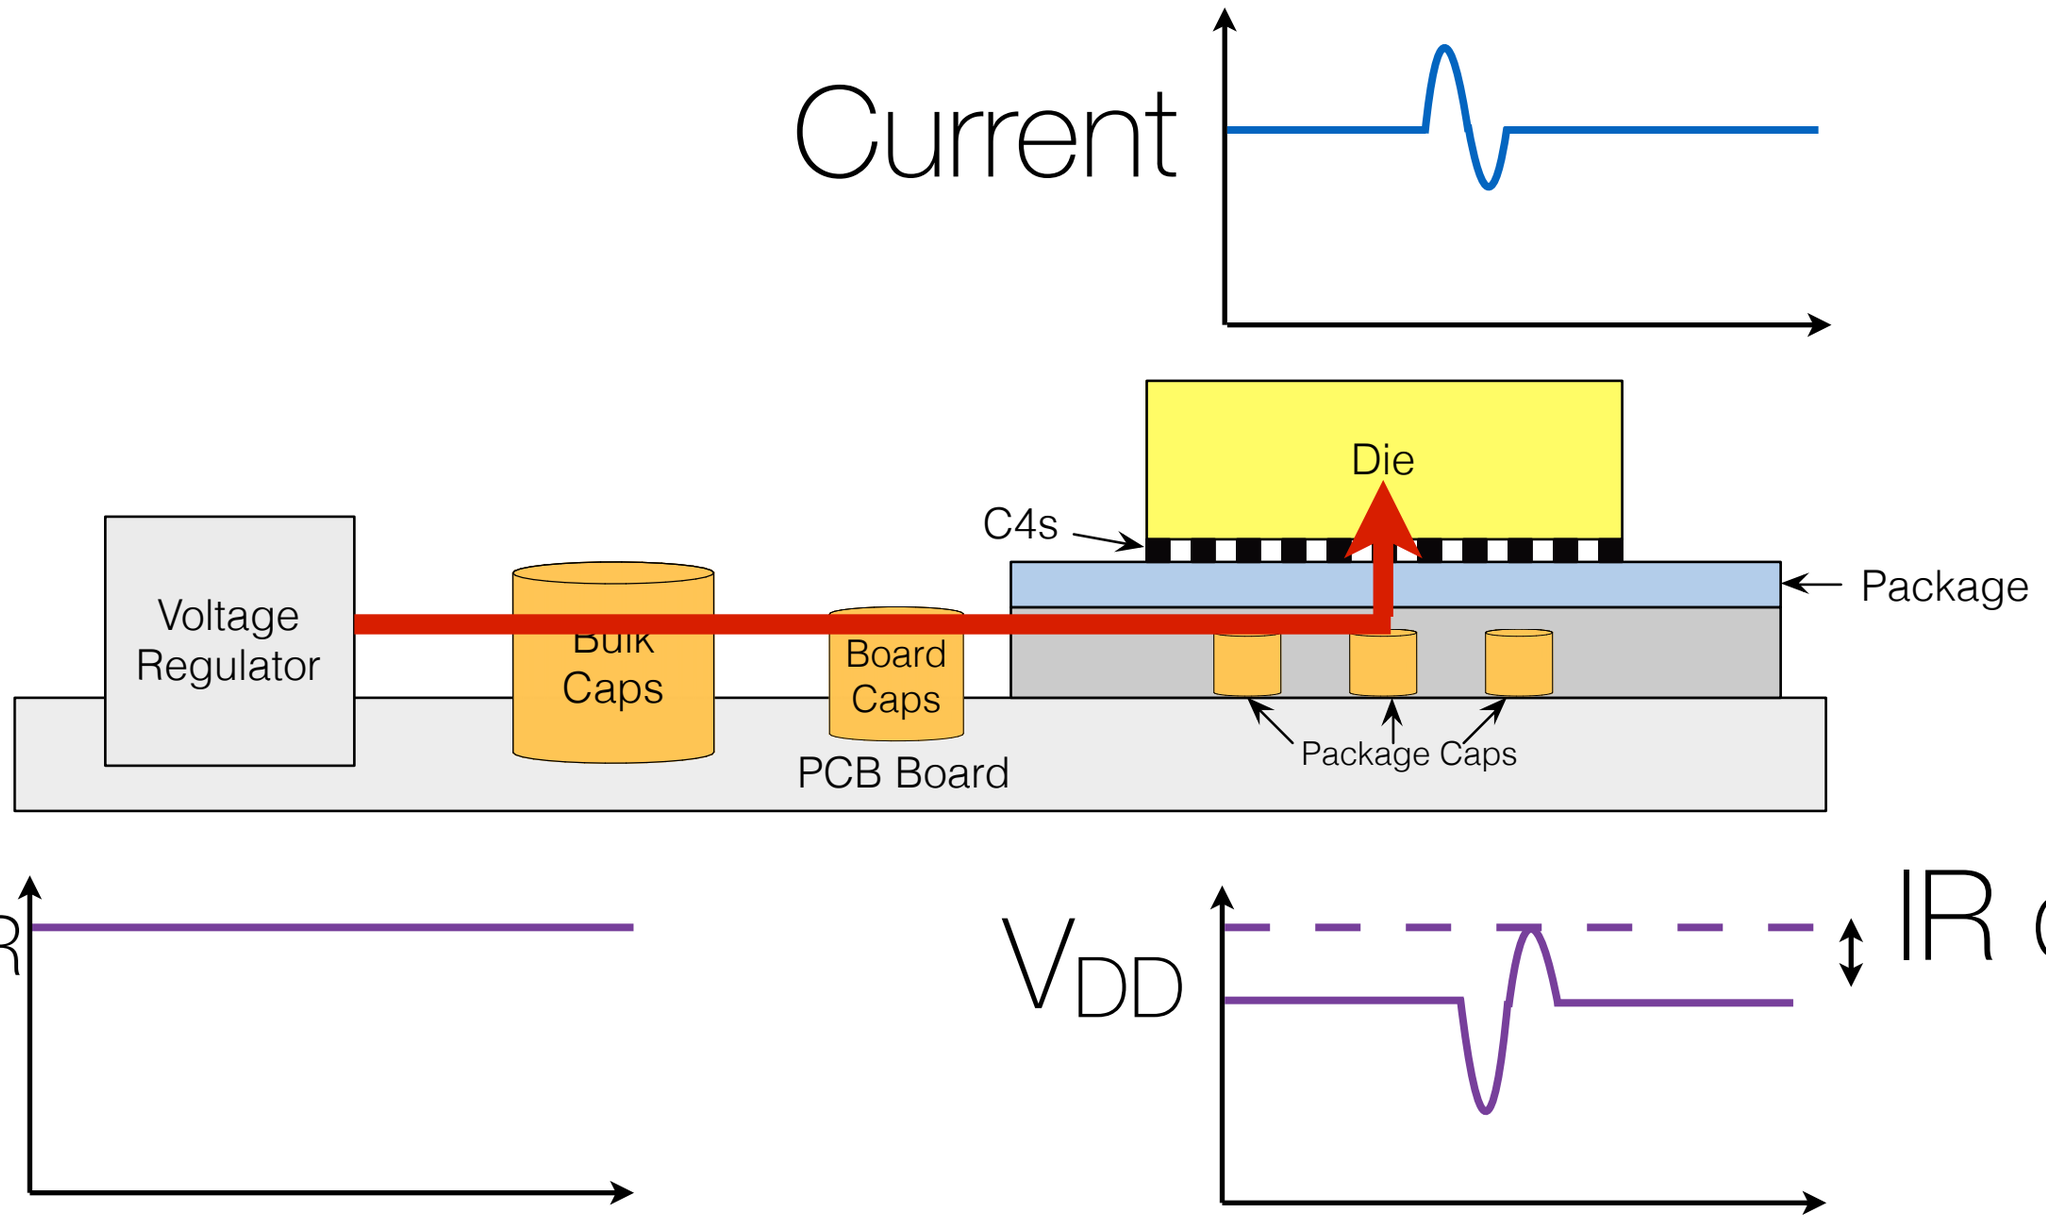

Voltage Noise Background

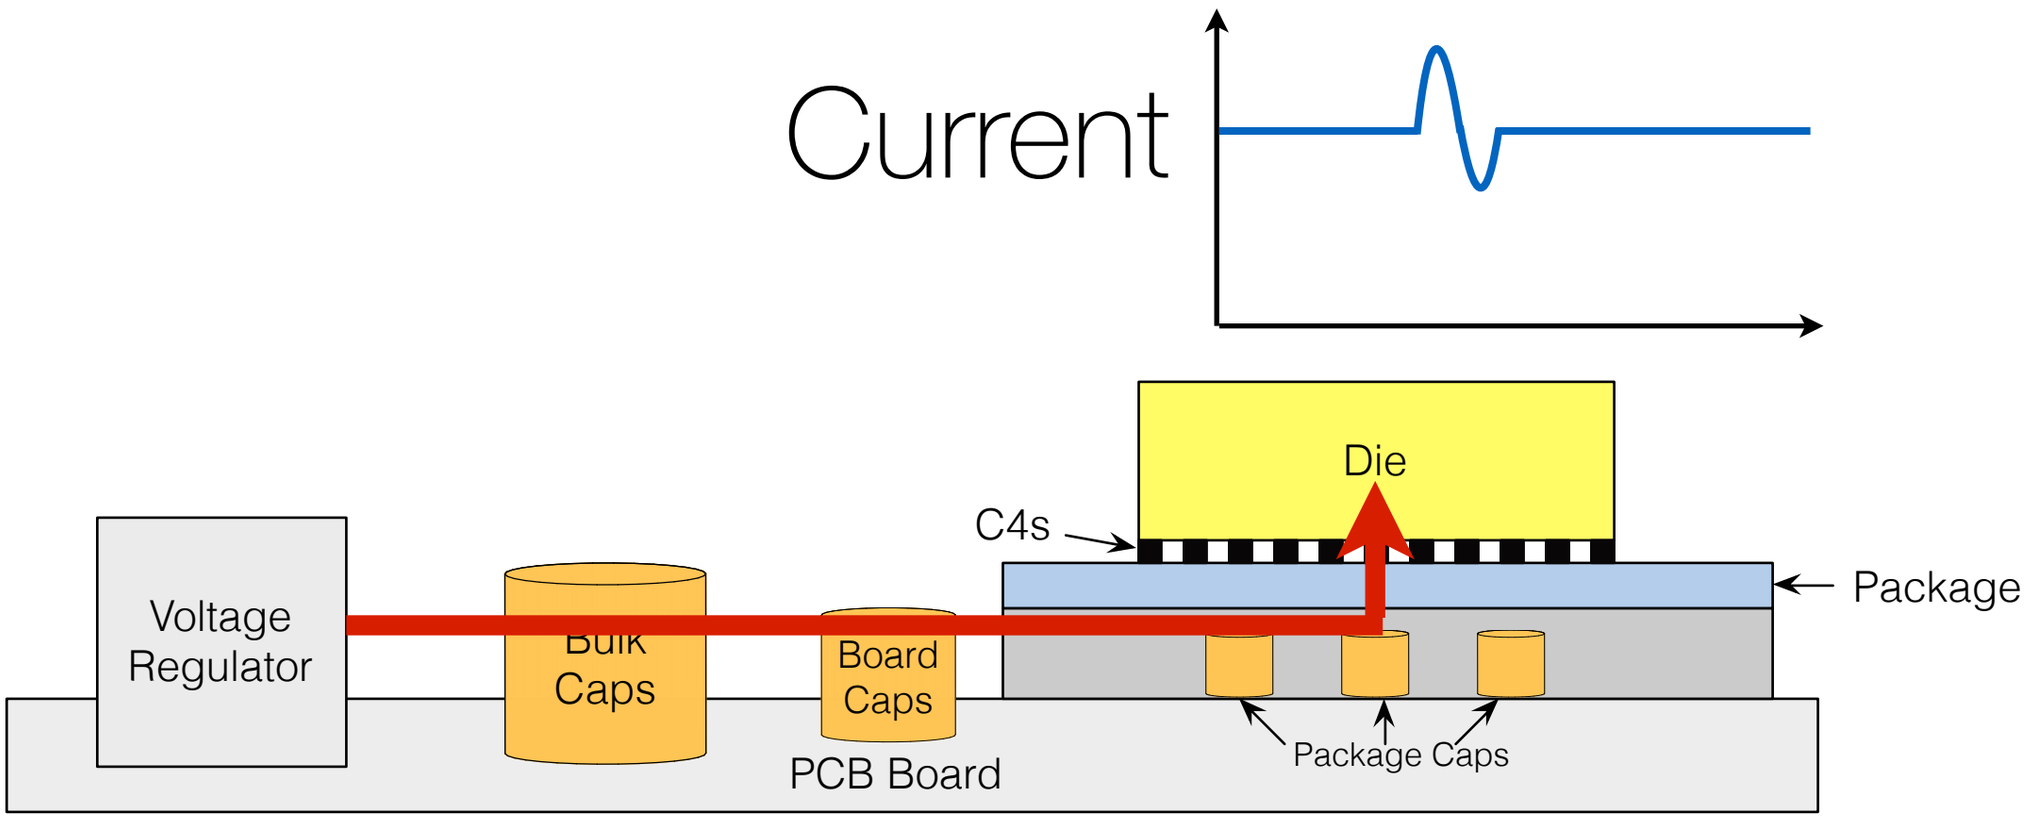

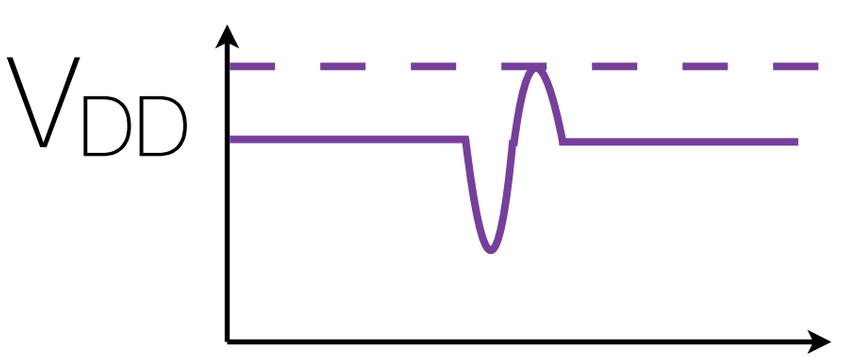

IR drop

di/dt

droop

Where does di/dt droop come from?

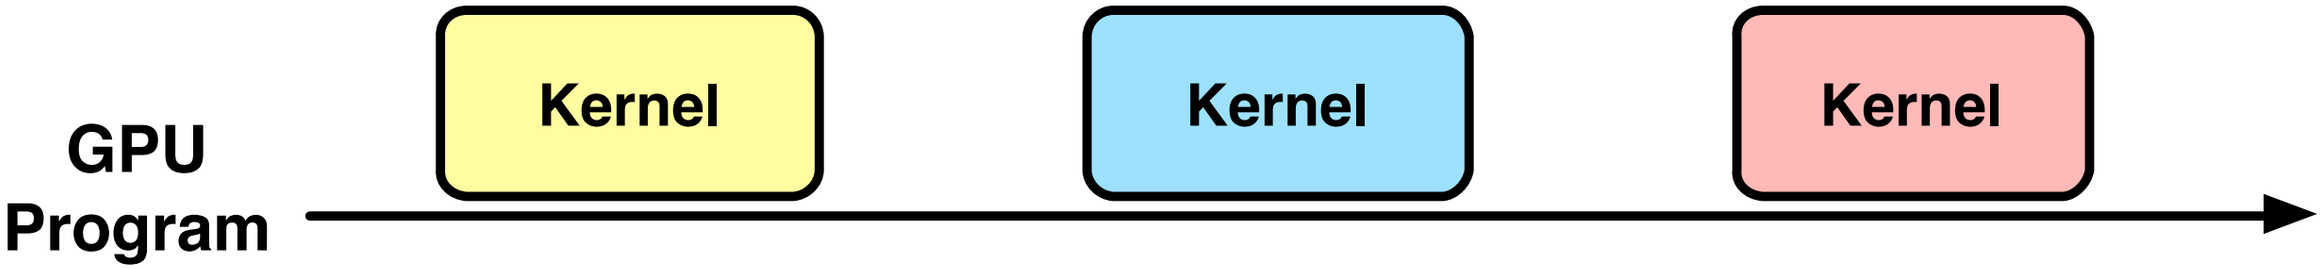

Where does di/dt droop come from?

- Kernel based activity patterns

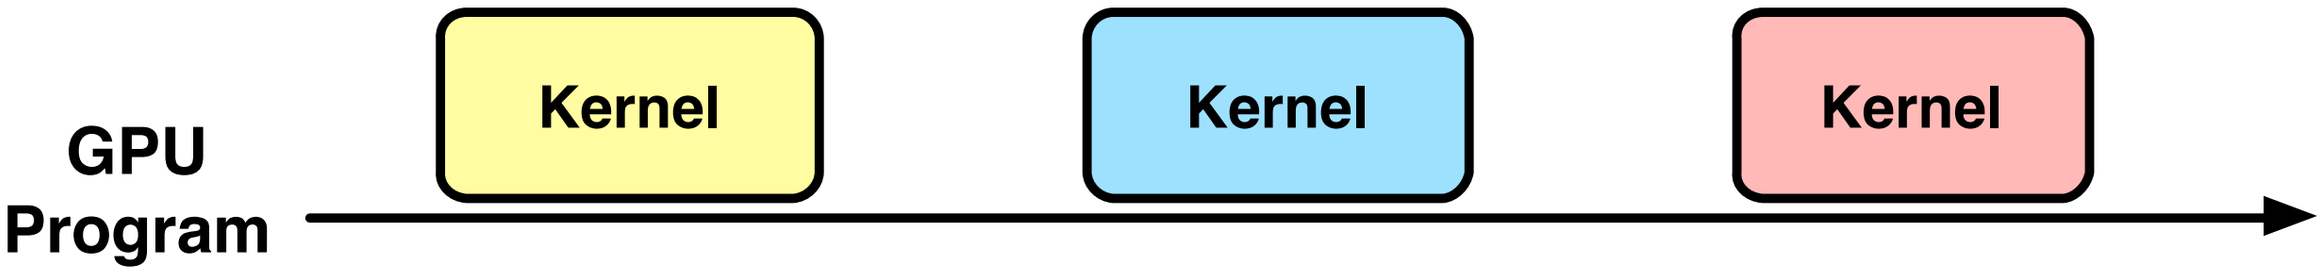

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

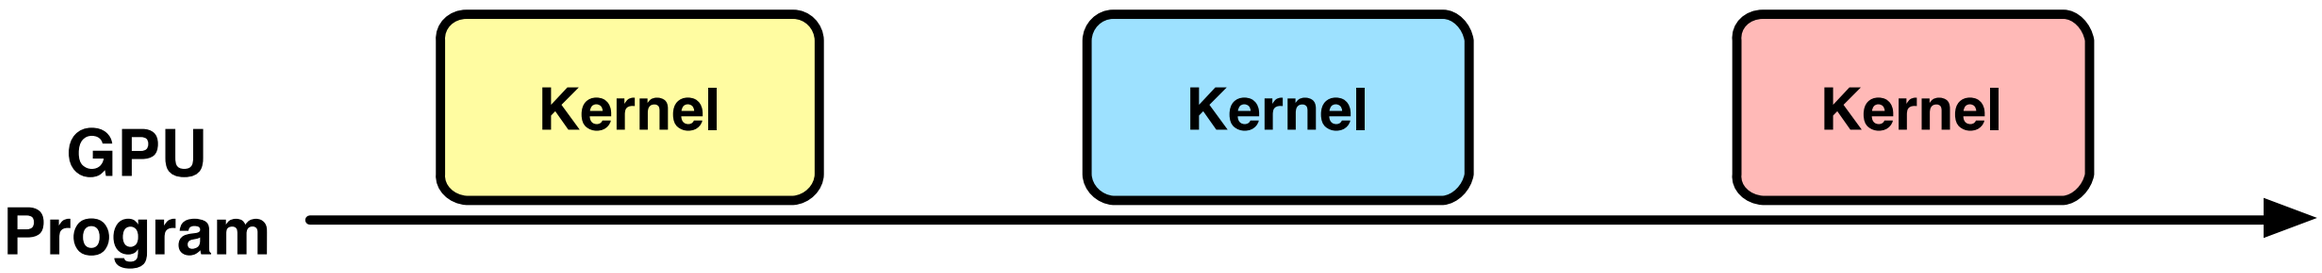

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

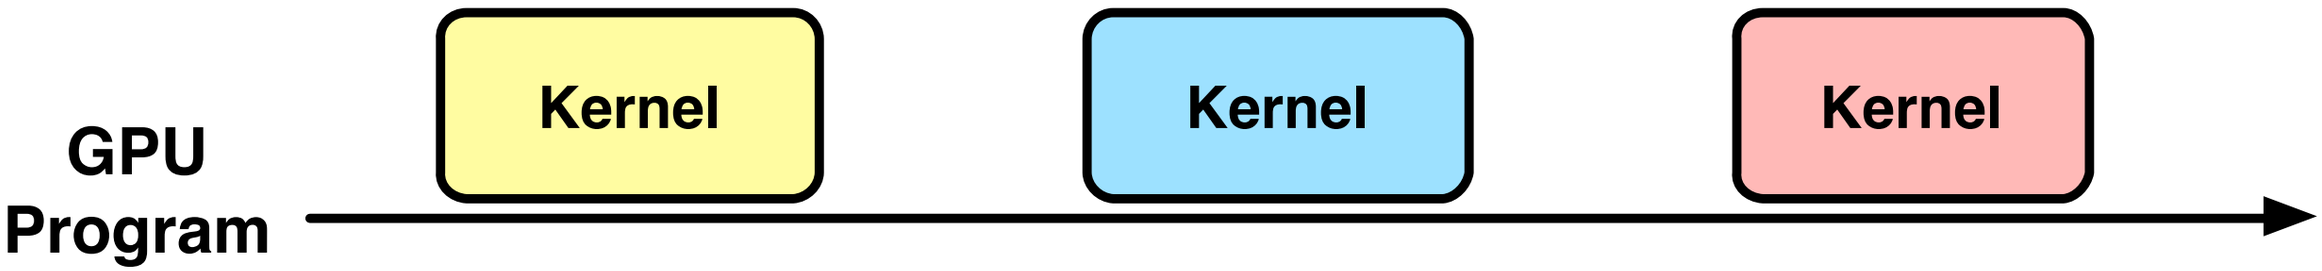

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

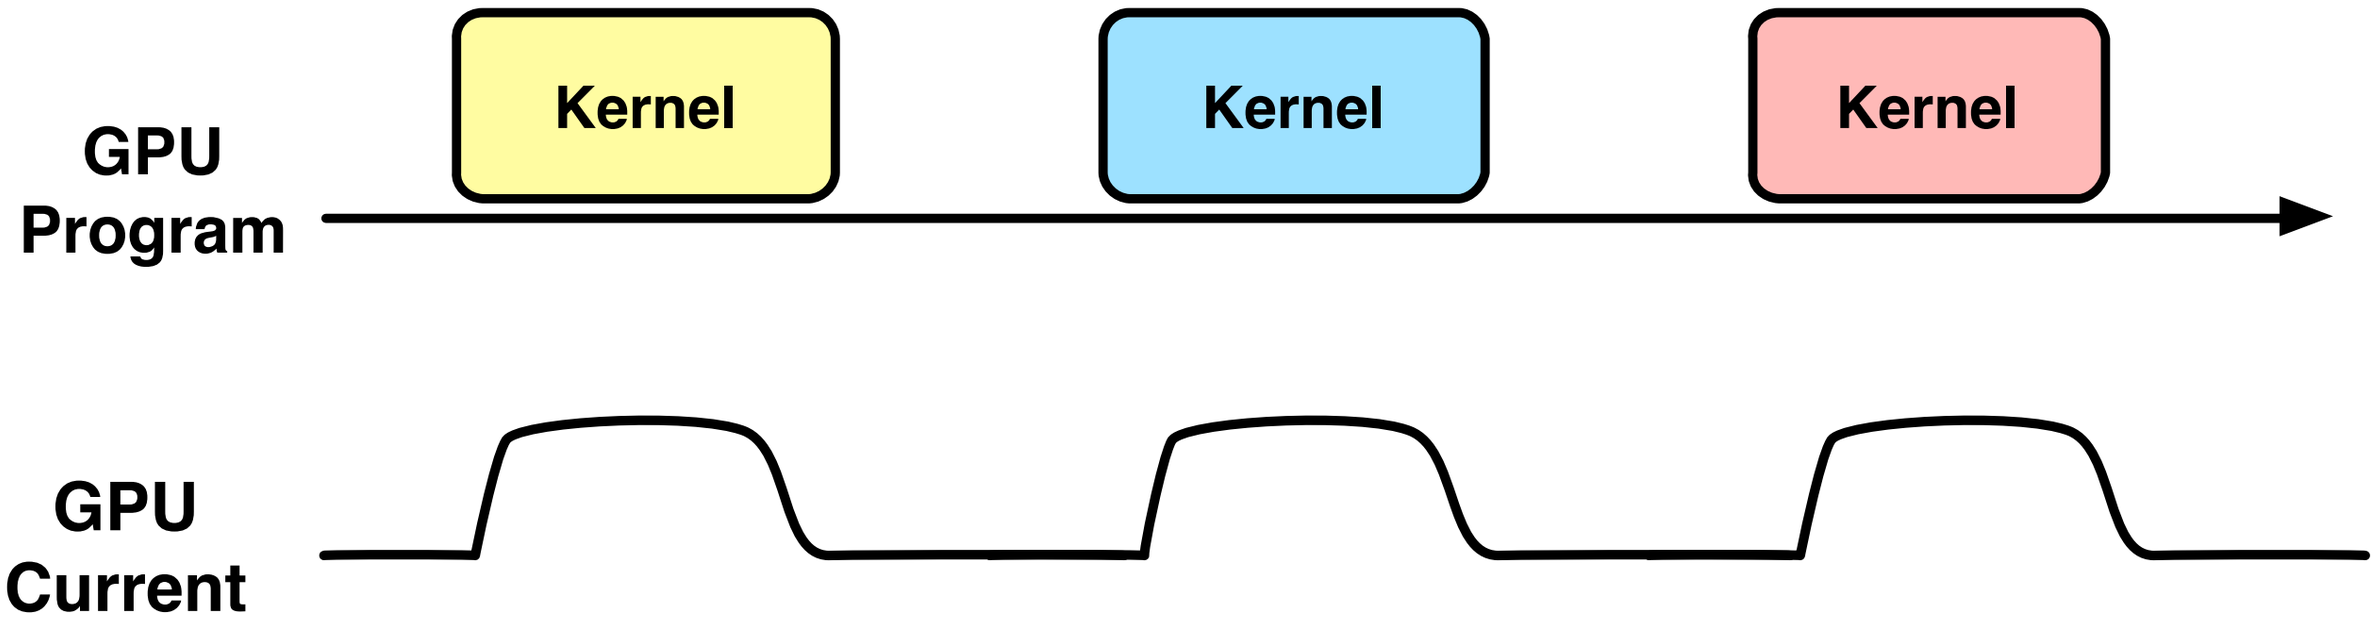

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

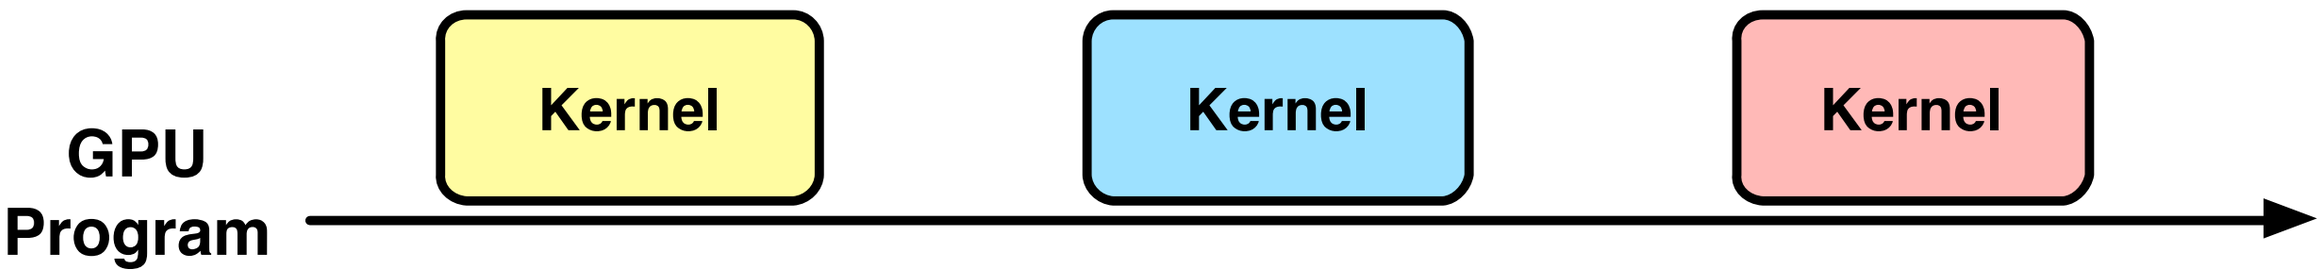

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

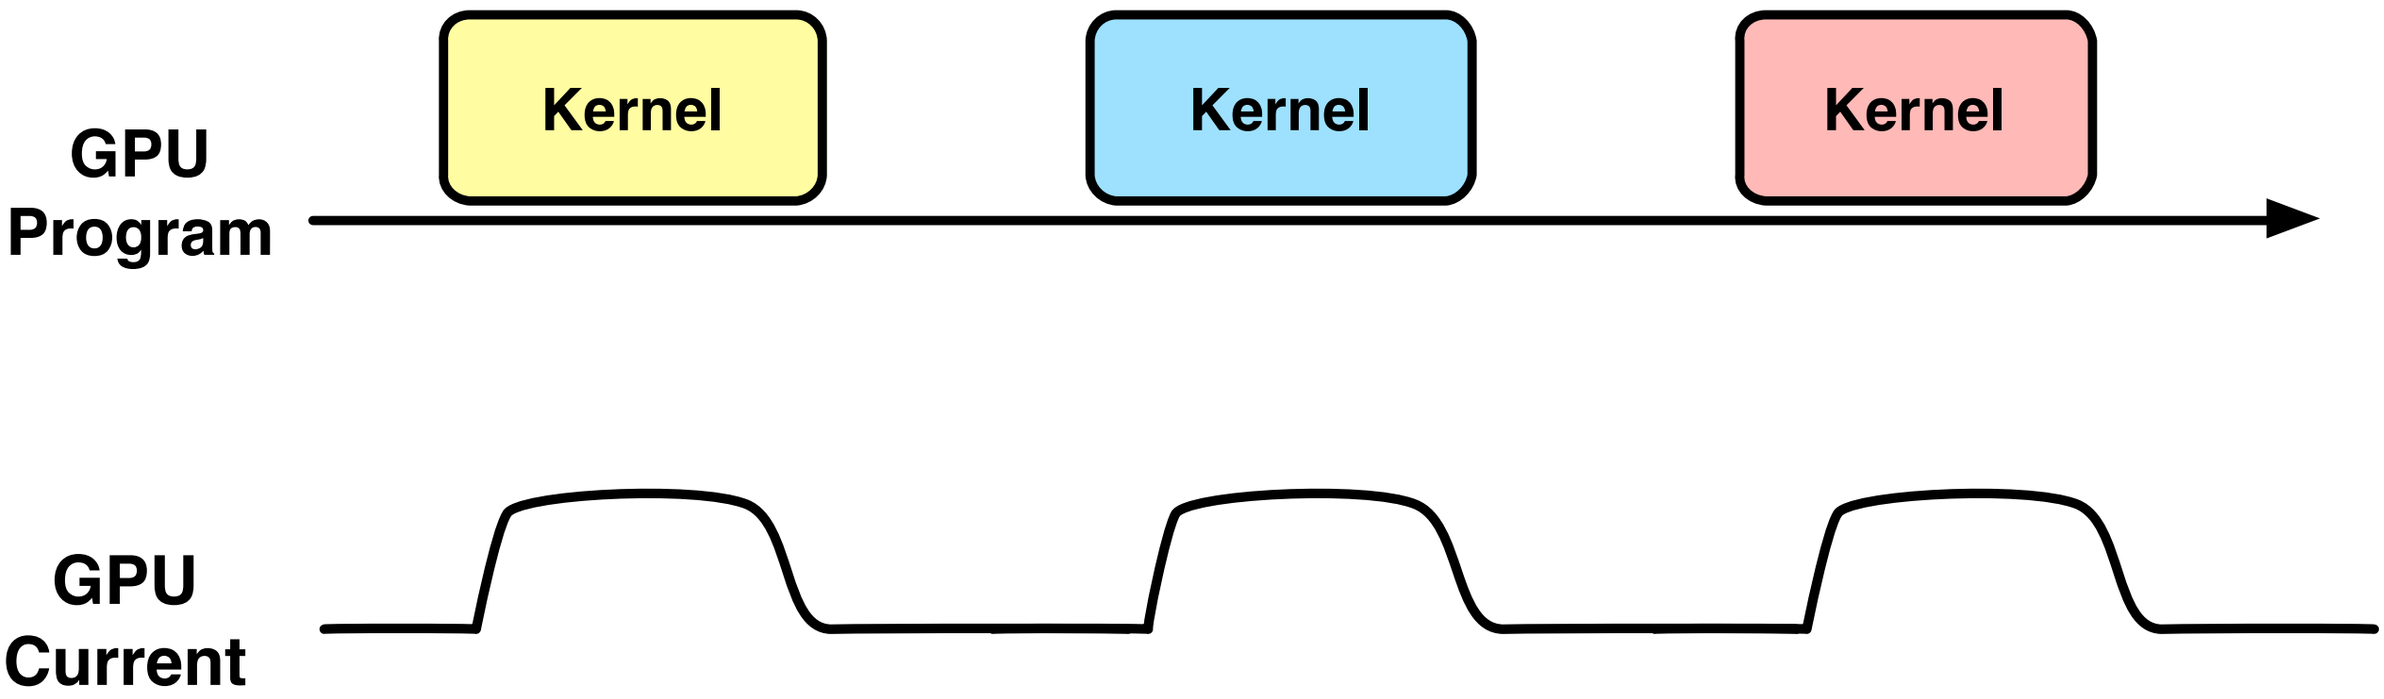

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

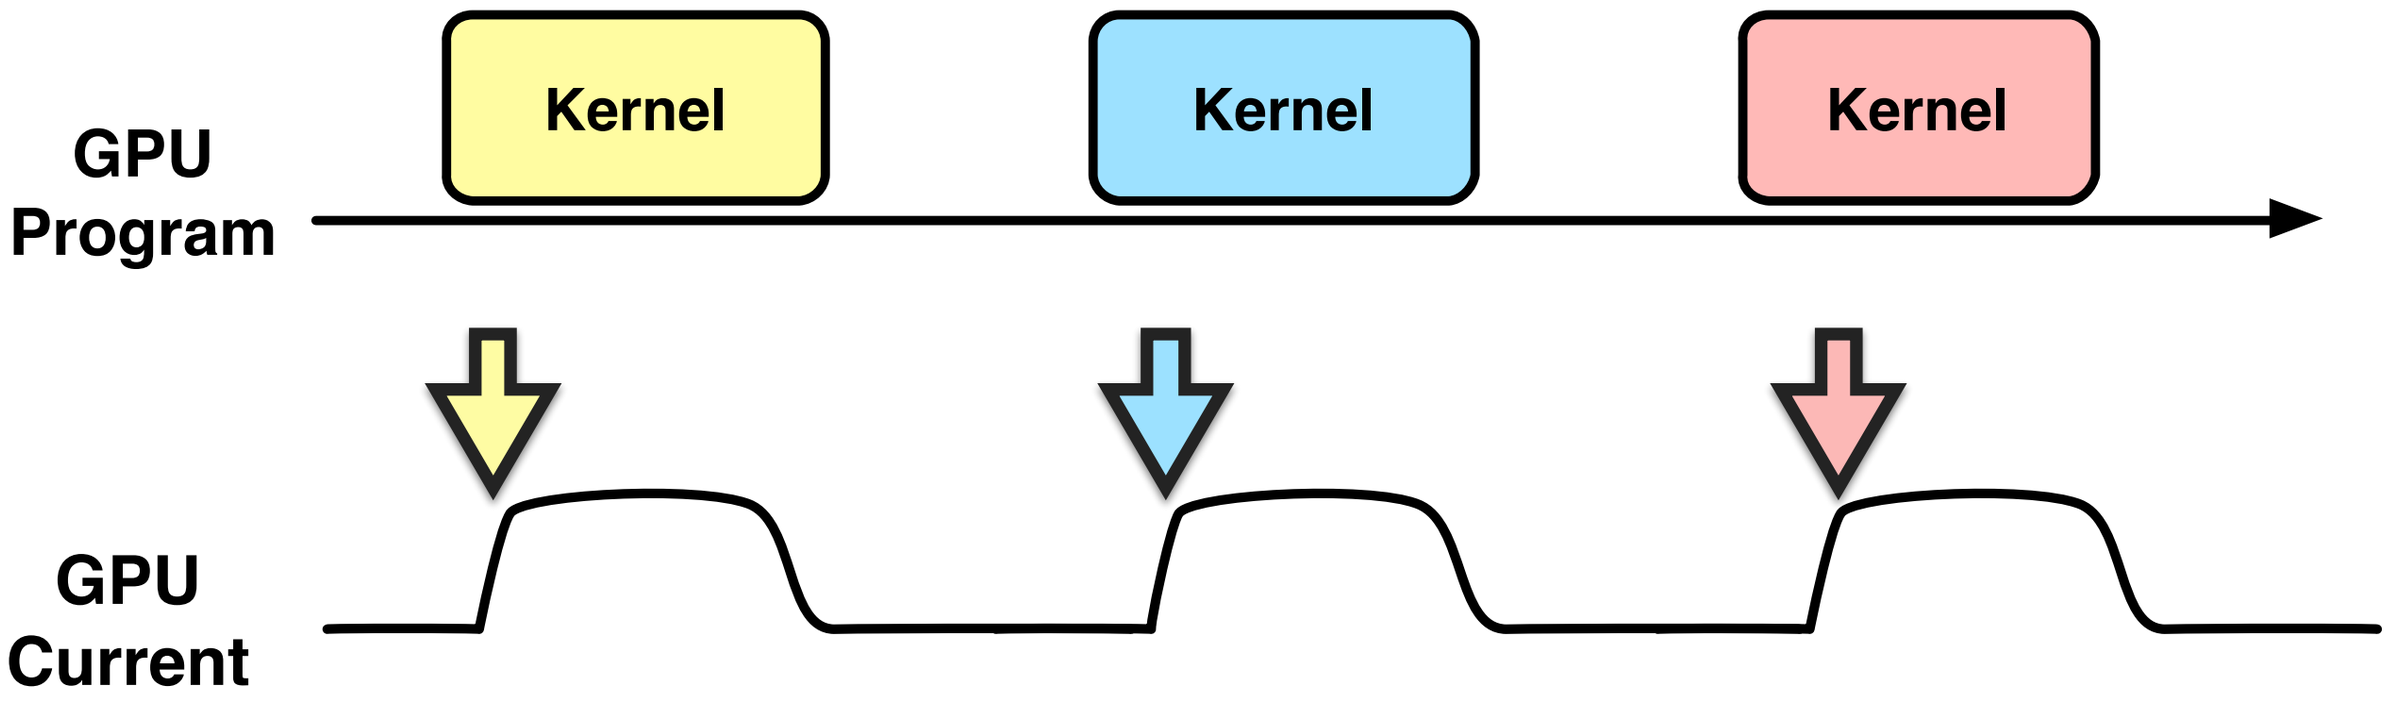

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

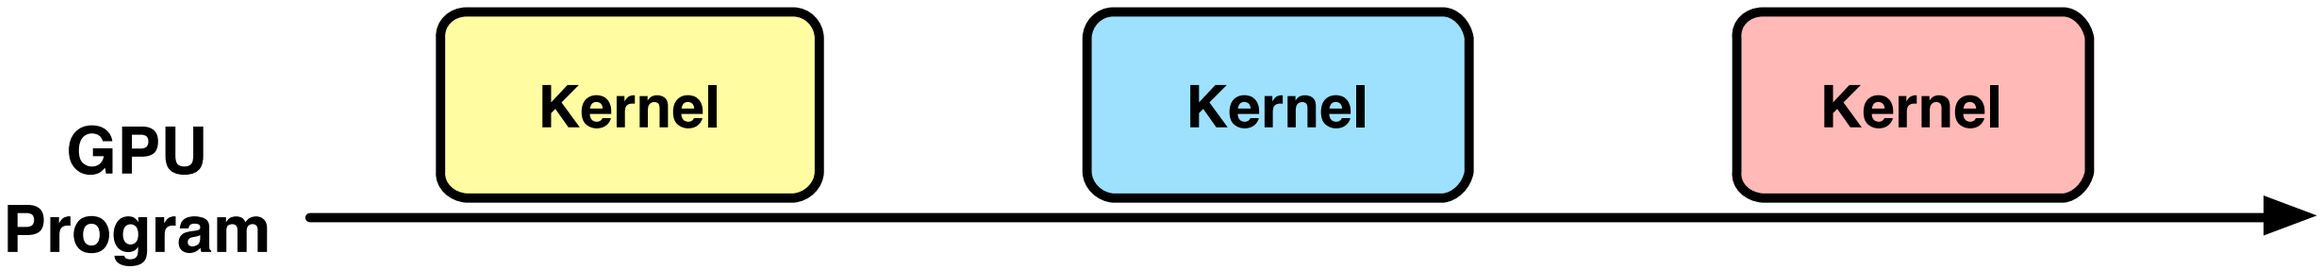

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

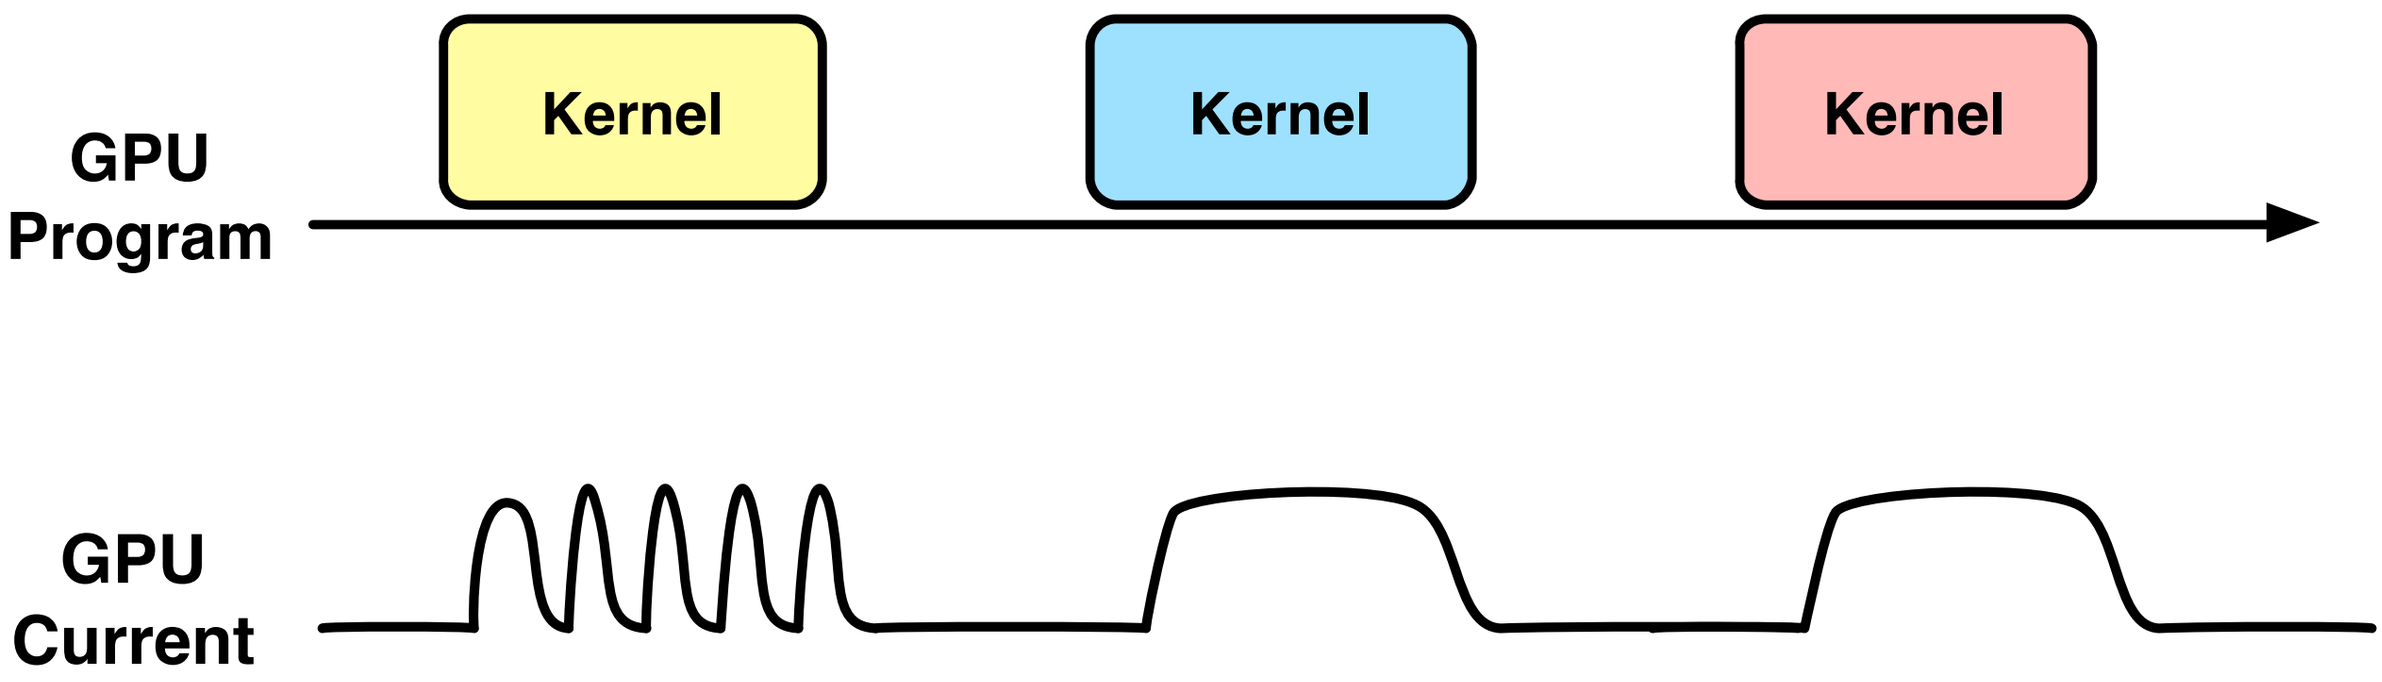

Where does di/dt droop come from?

- Kernel based activity patterns
  - Inter kernel
  - Initial kernel
  - Intra kernel

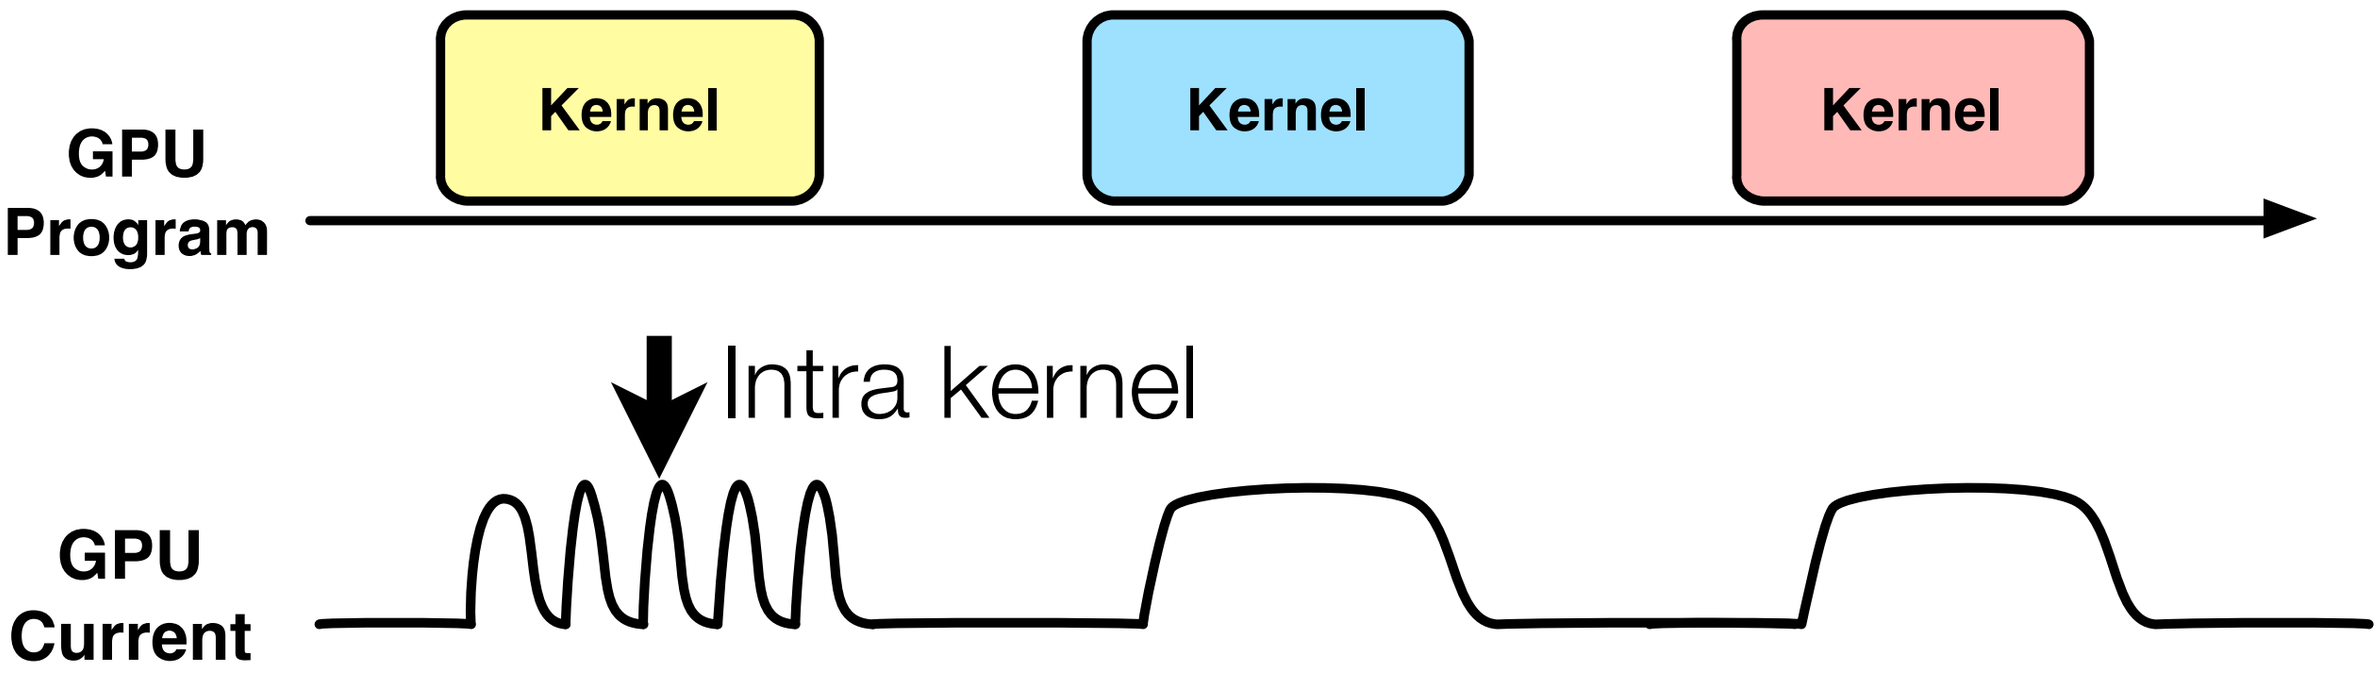

Kernel Level V_{min} Measurement

Kernel Level V_{min} Measurement

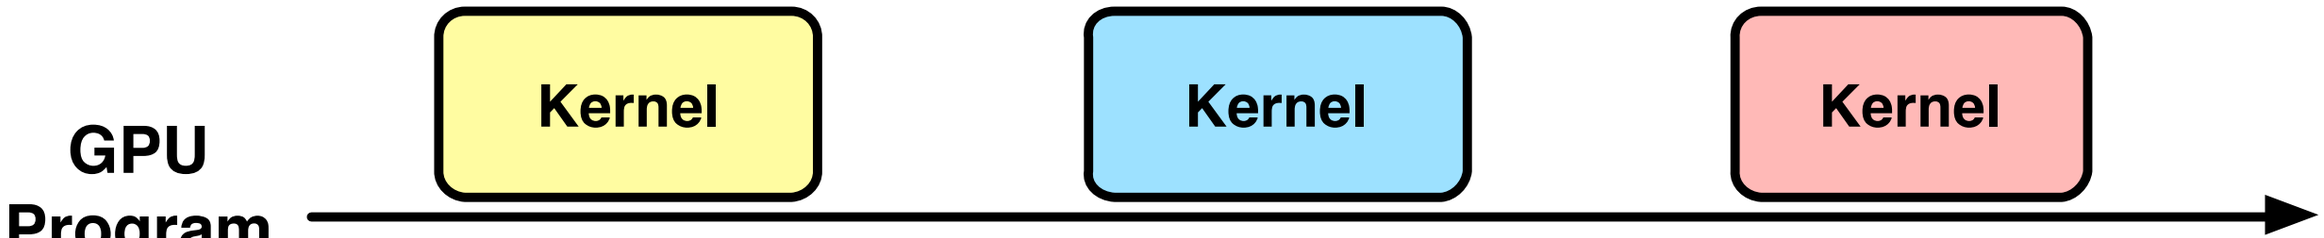

Kernel Level V_{min} Measurement

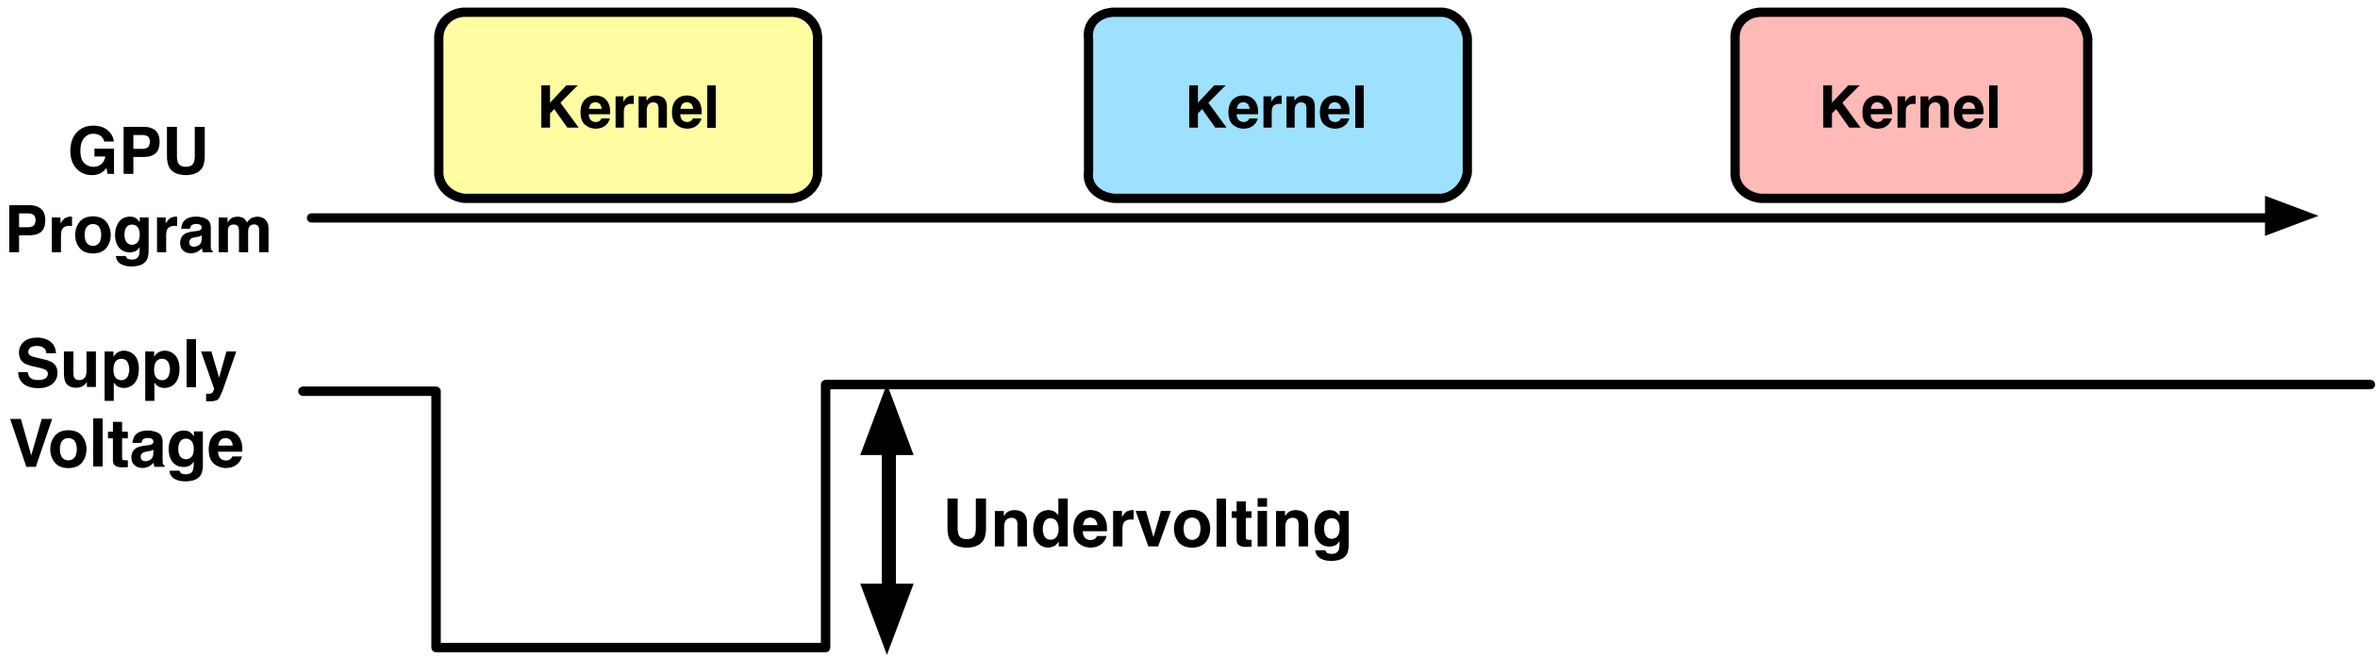

Kernel Level V_{min} Measurement

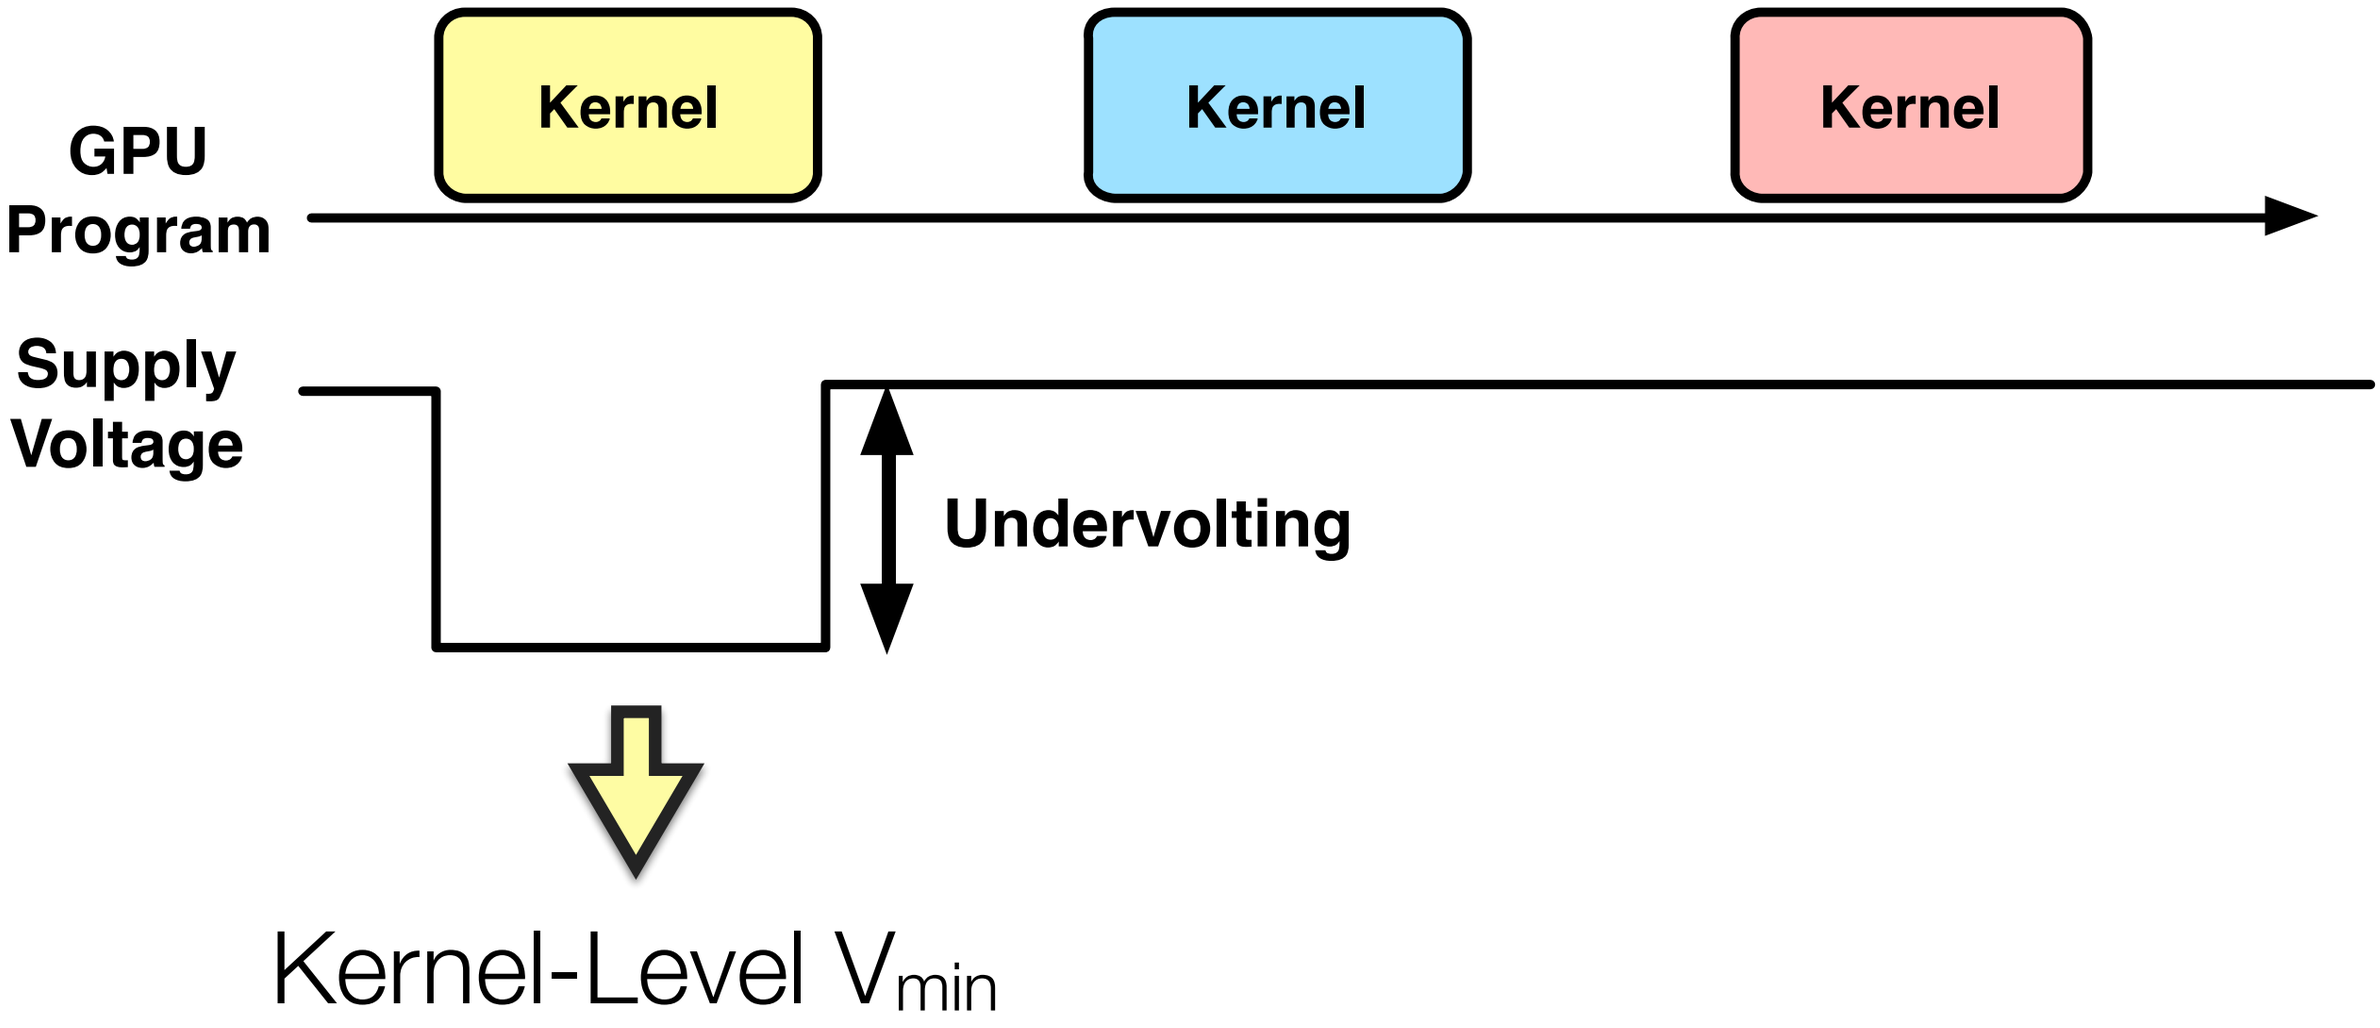

Kernel Level V_{min} Measurement

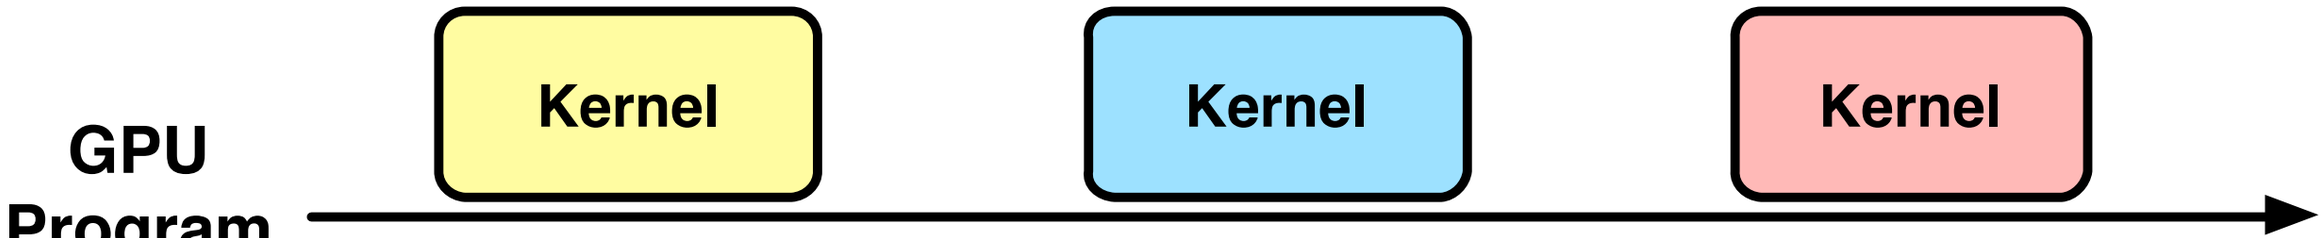

Kernel Level V_{min} Measurement

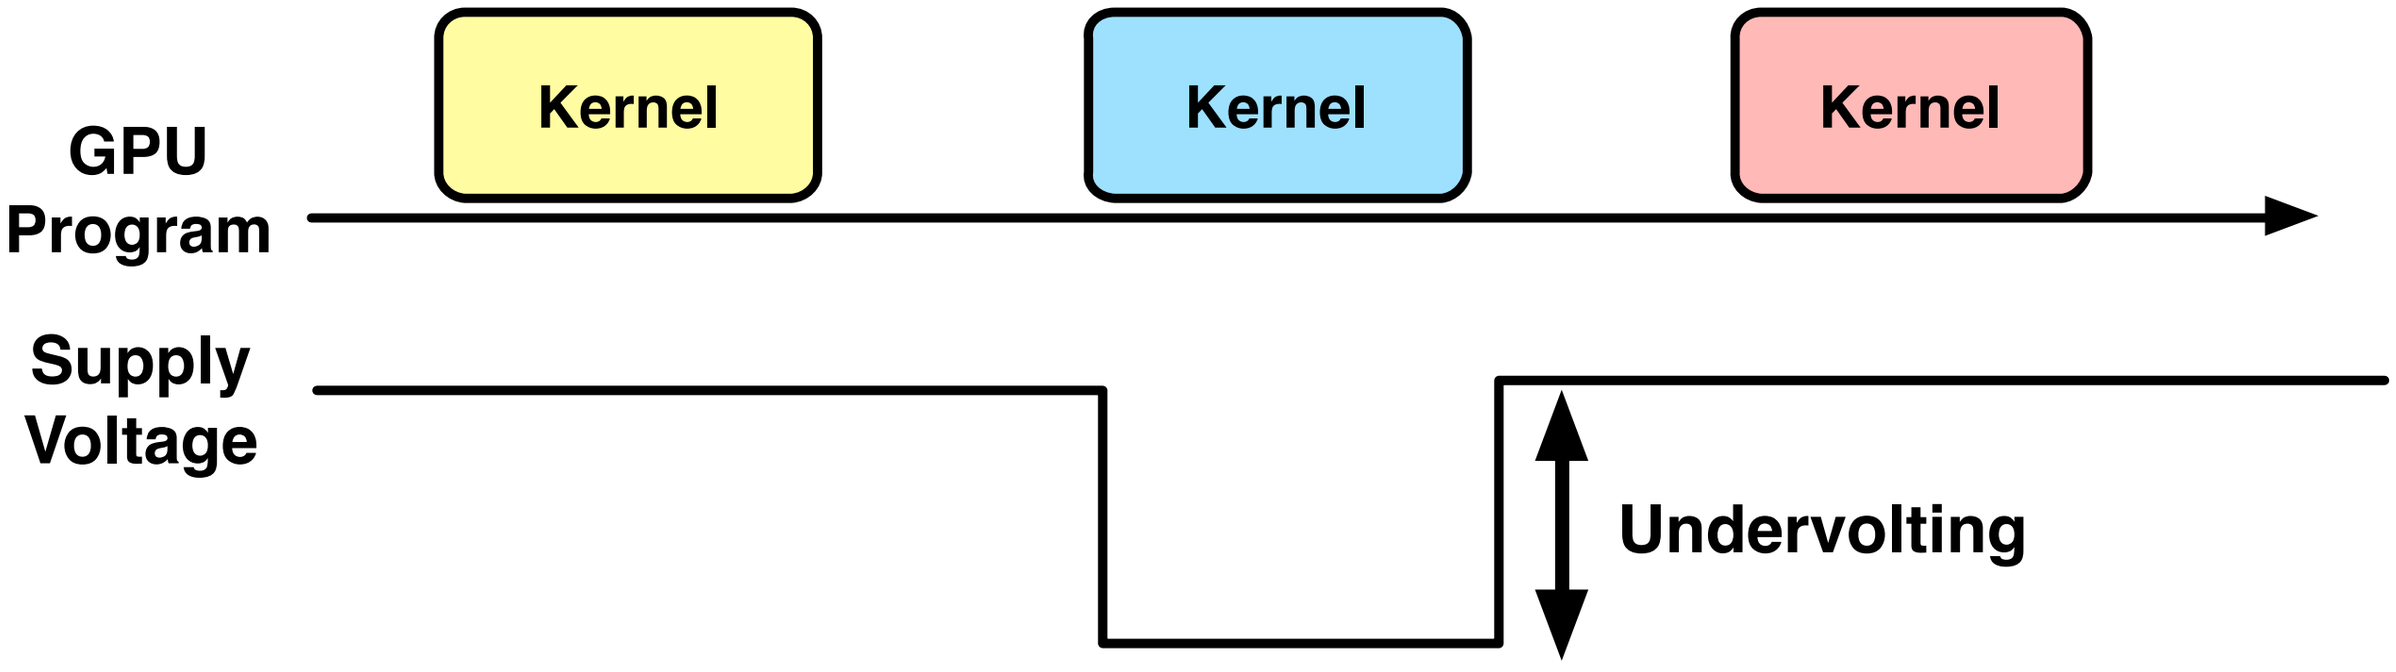

Kernel Level V_{min} Measurement

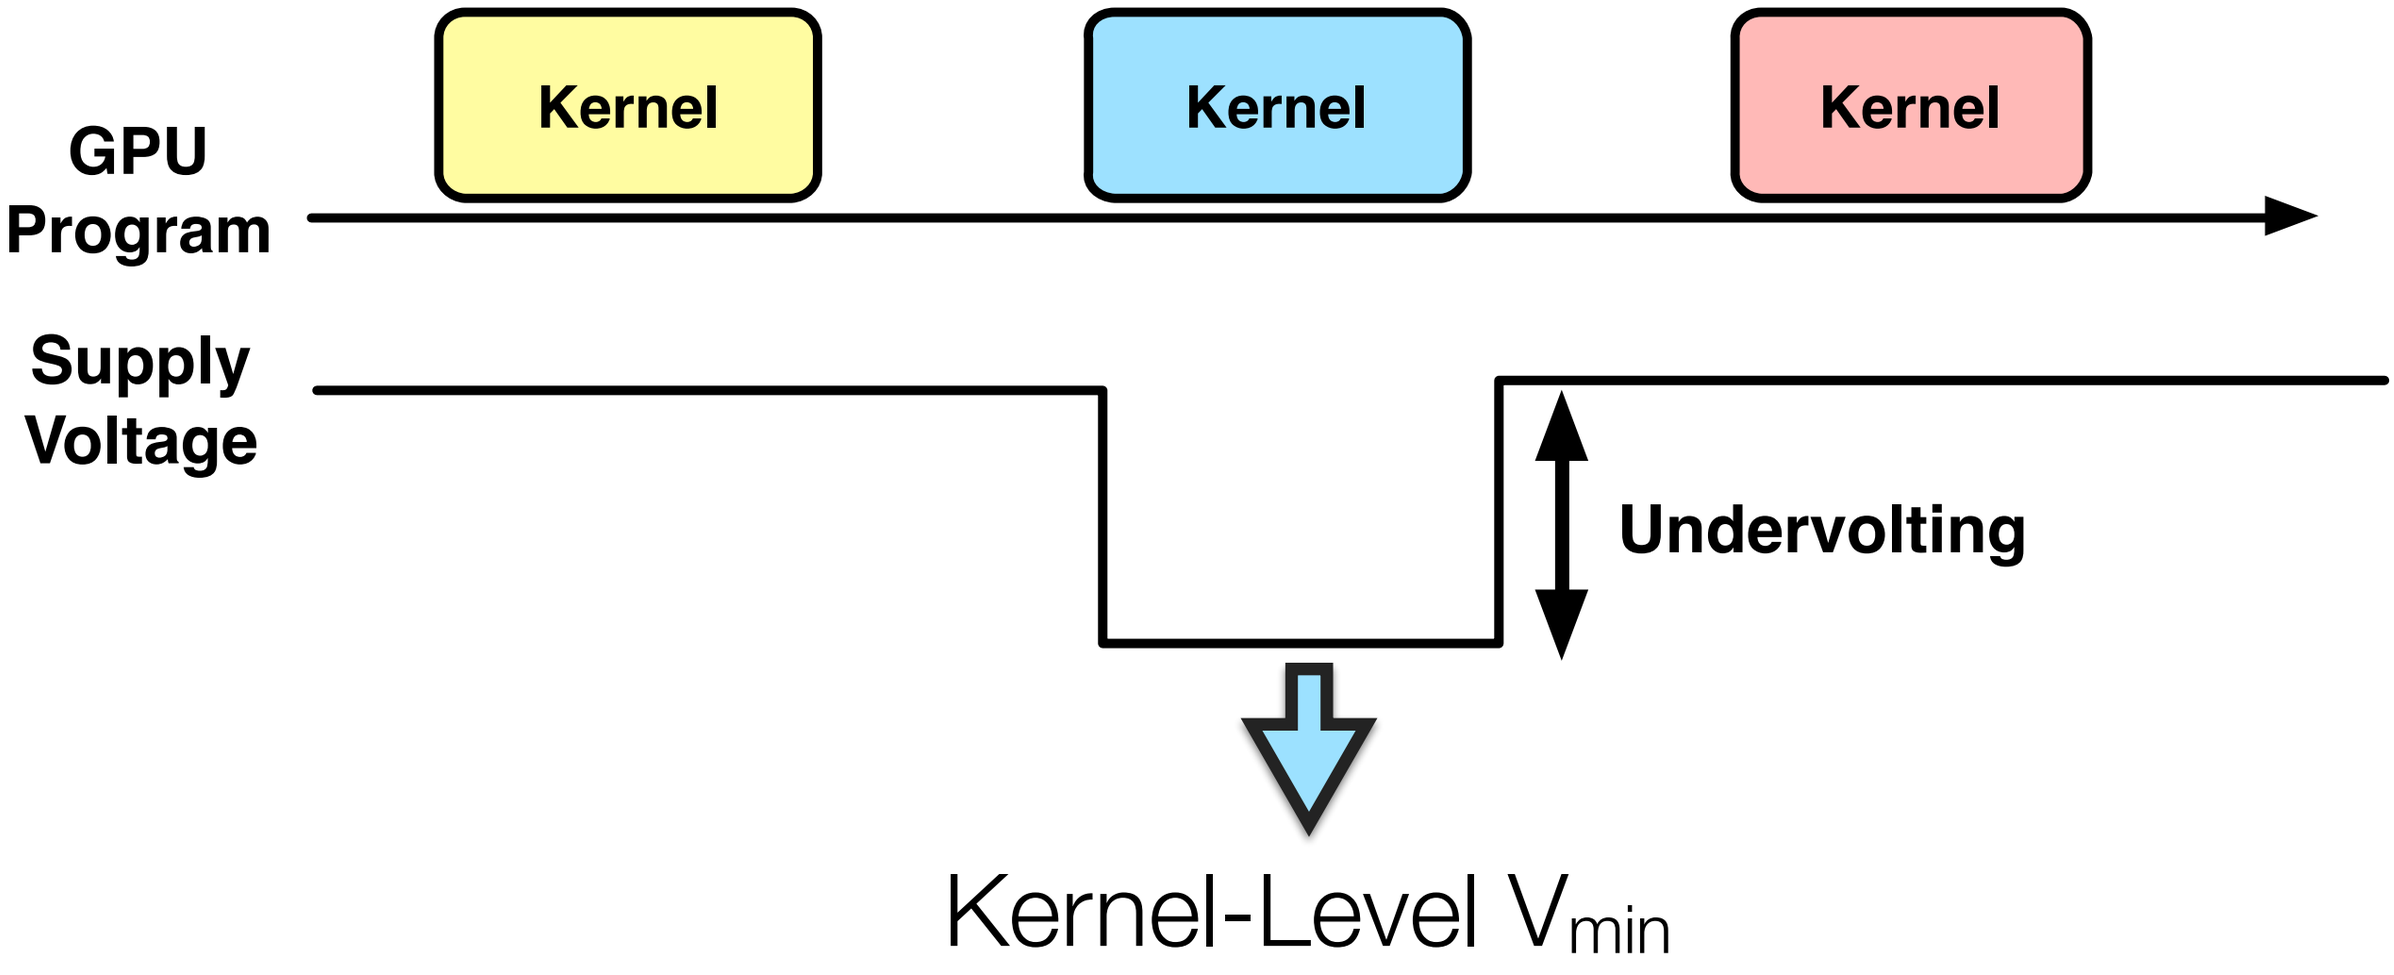

Program Level V_{min} Measurement

Program Level V_{min} Measurement

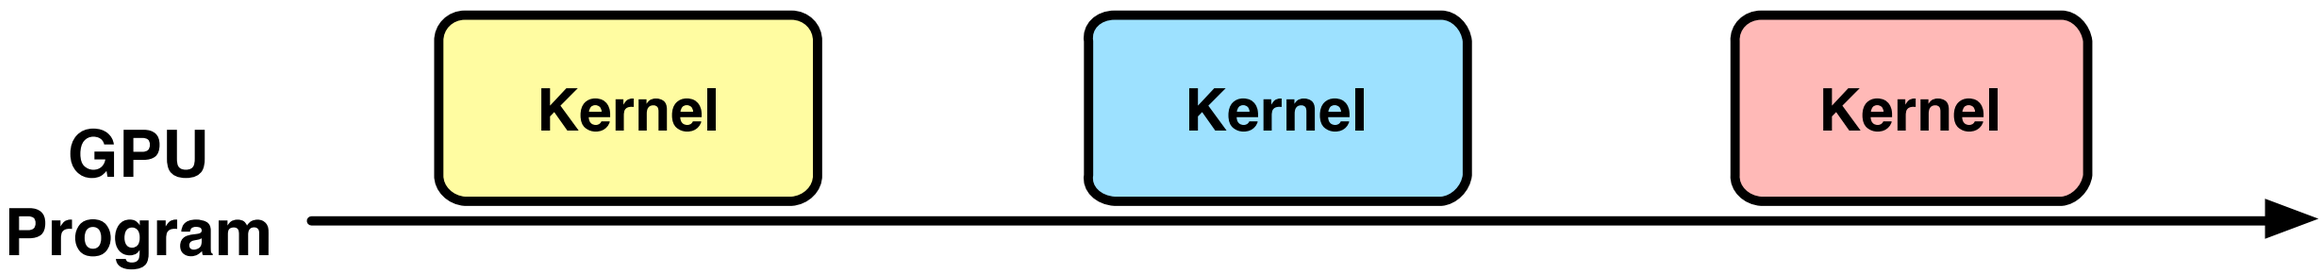

Program Level V_{min} Measurement

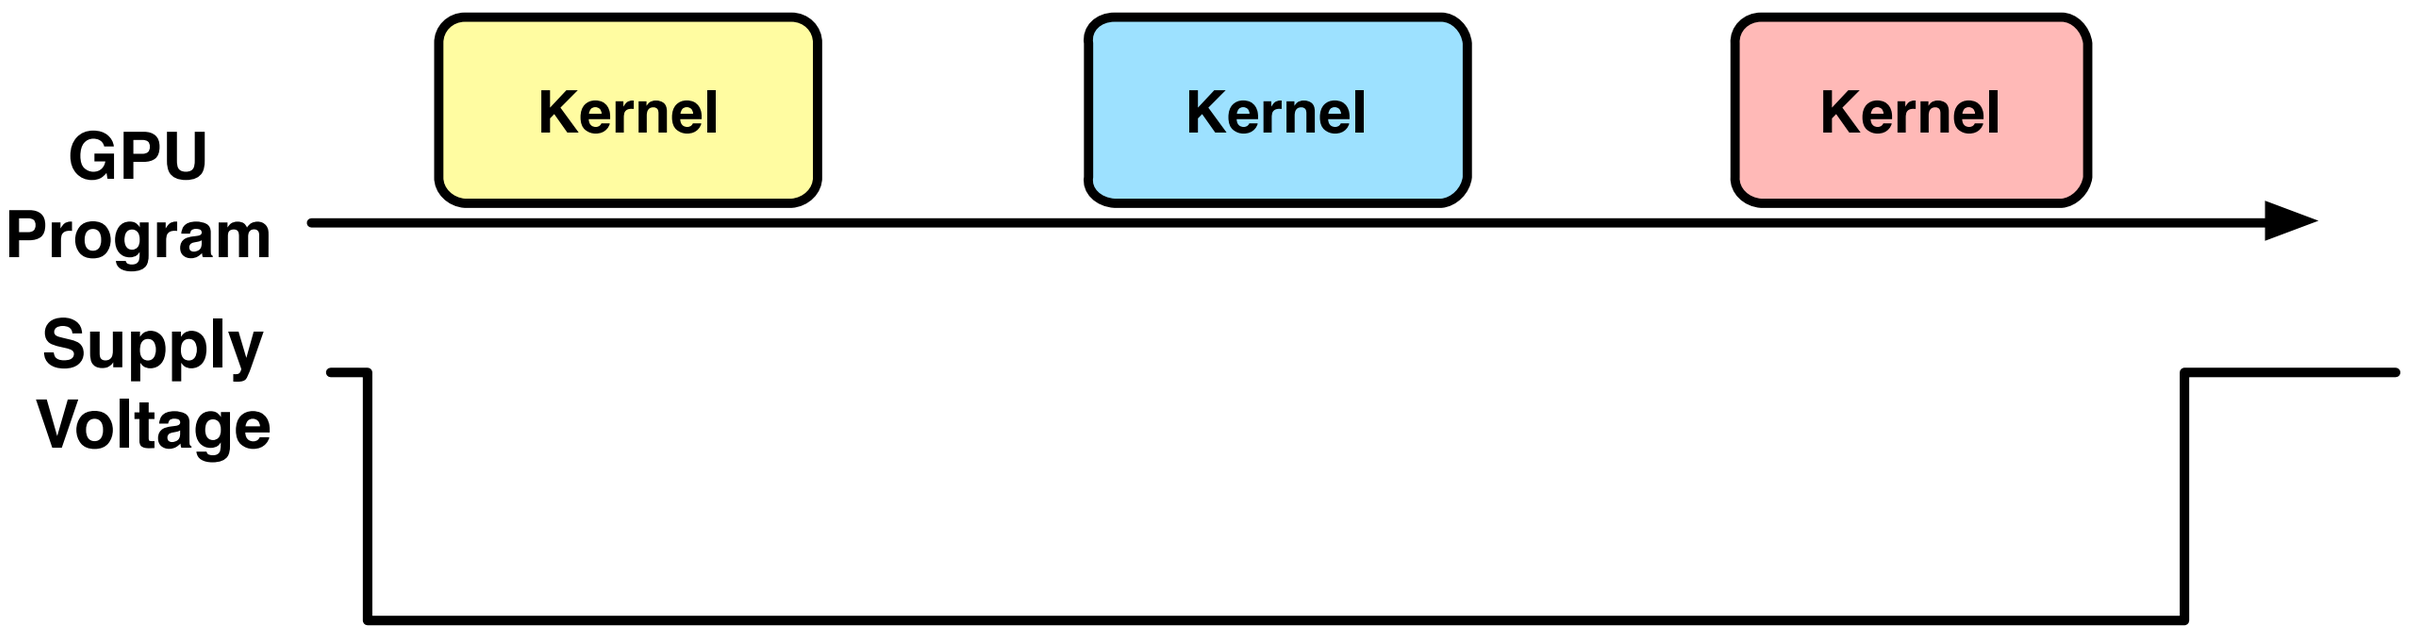

Program Level V_{min} Measurement

Program Level V_{min} Measurement

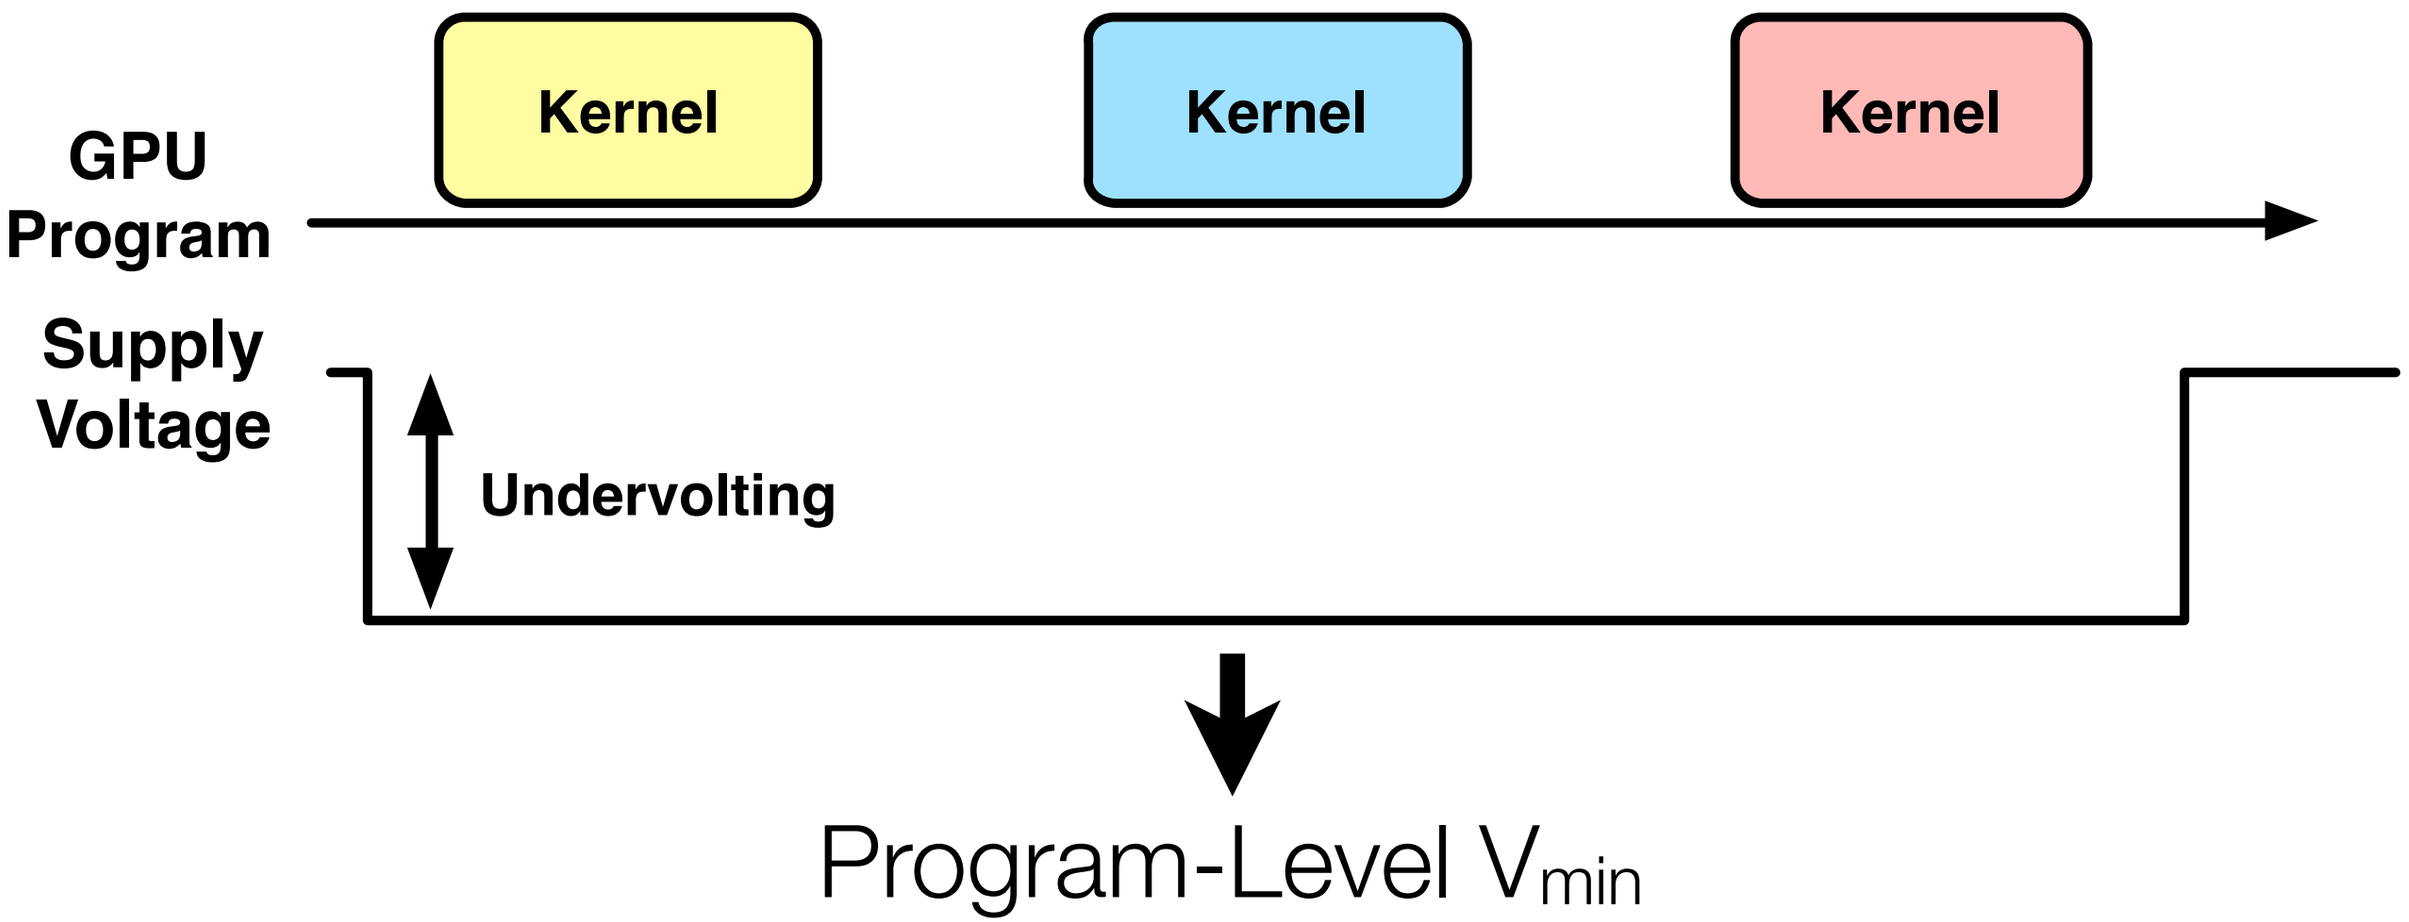

Program/Kernel Level V_{min} Comparison

Program/Kernel Level V_{min} Comparison

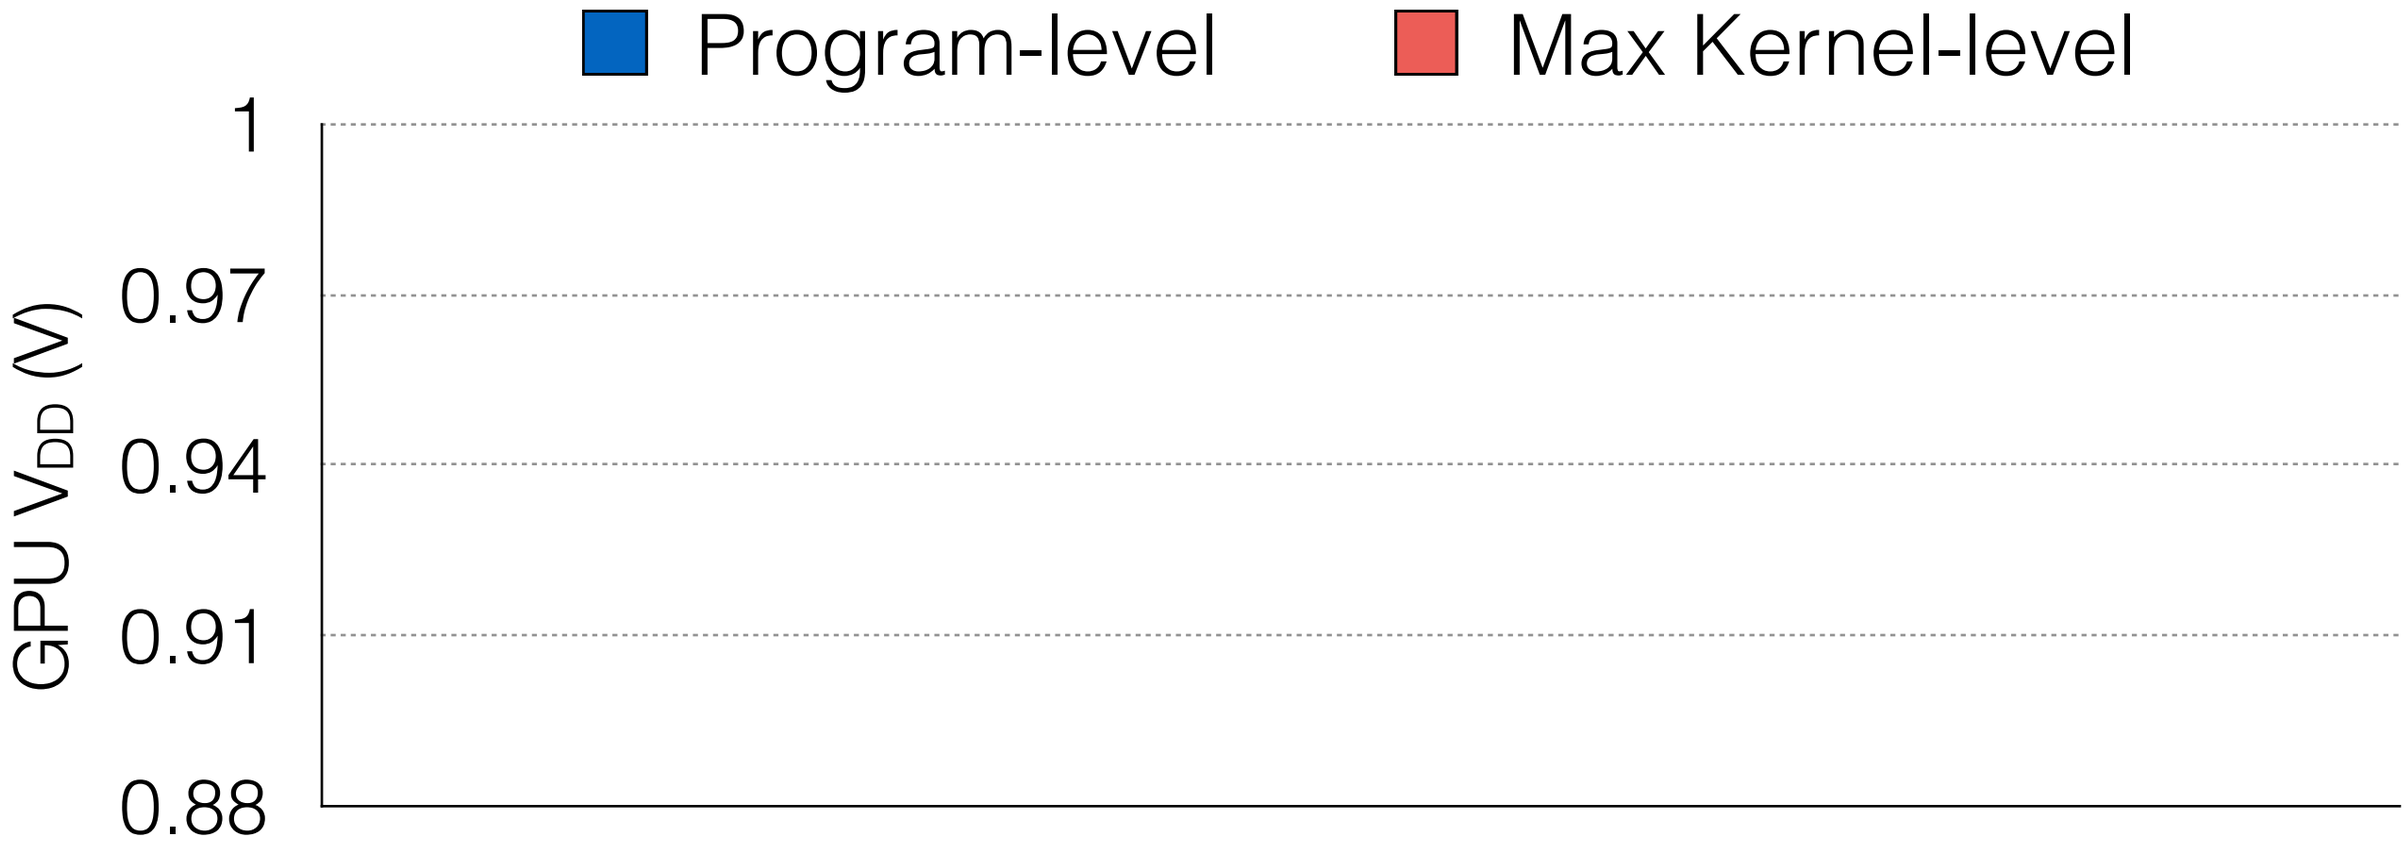

Program/Kernel Level V_{min} Comparison

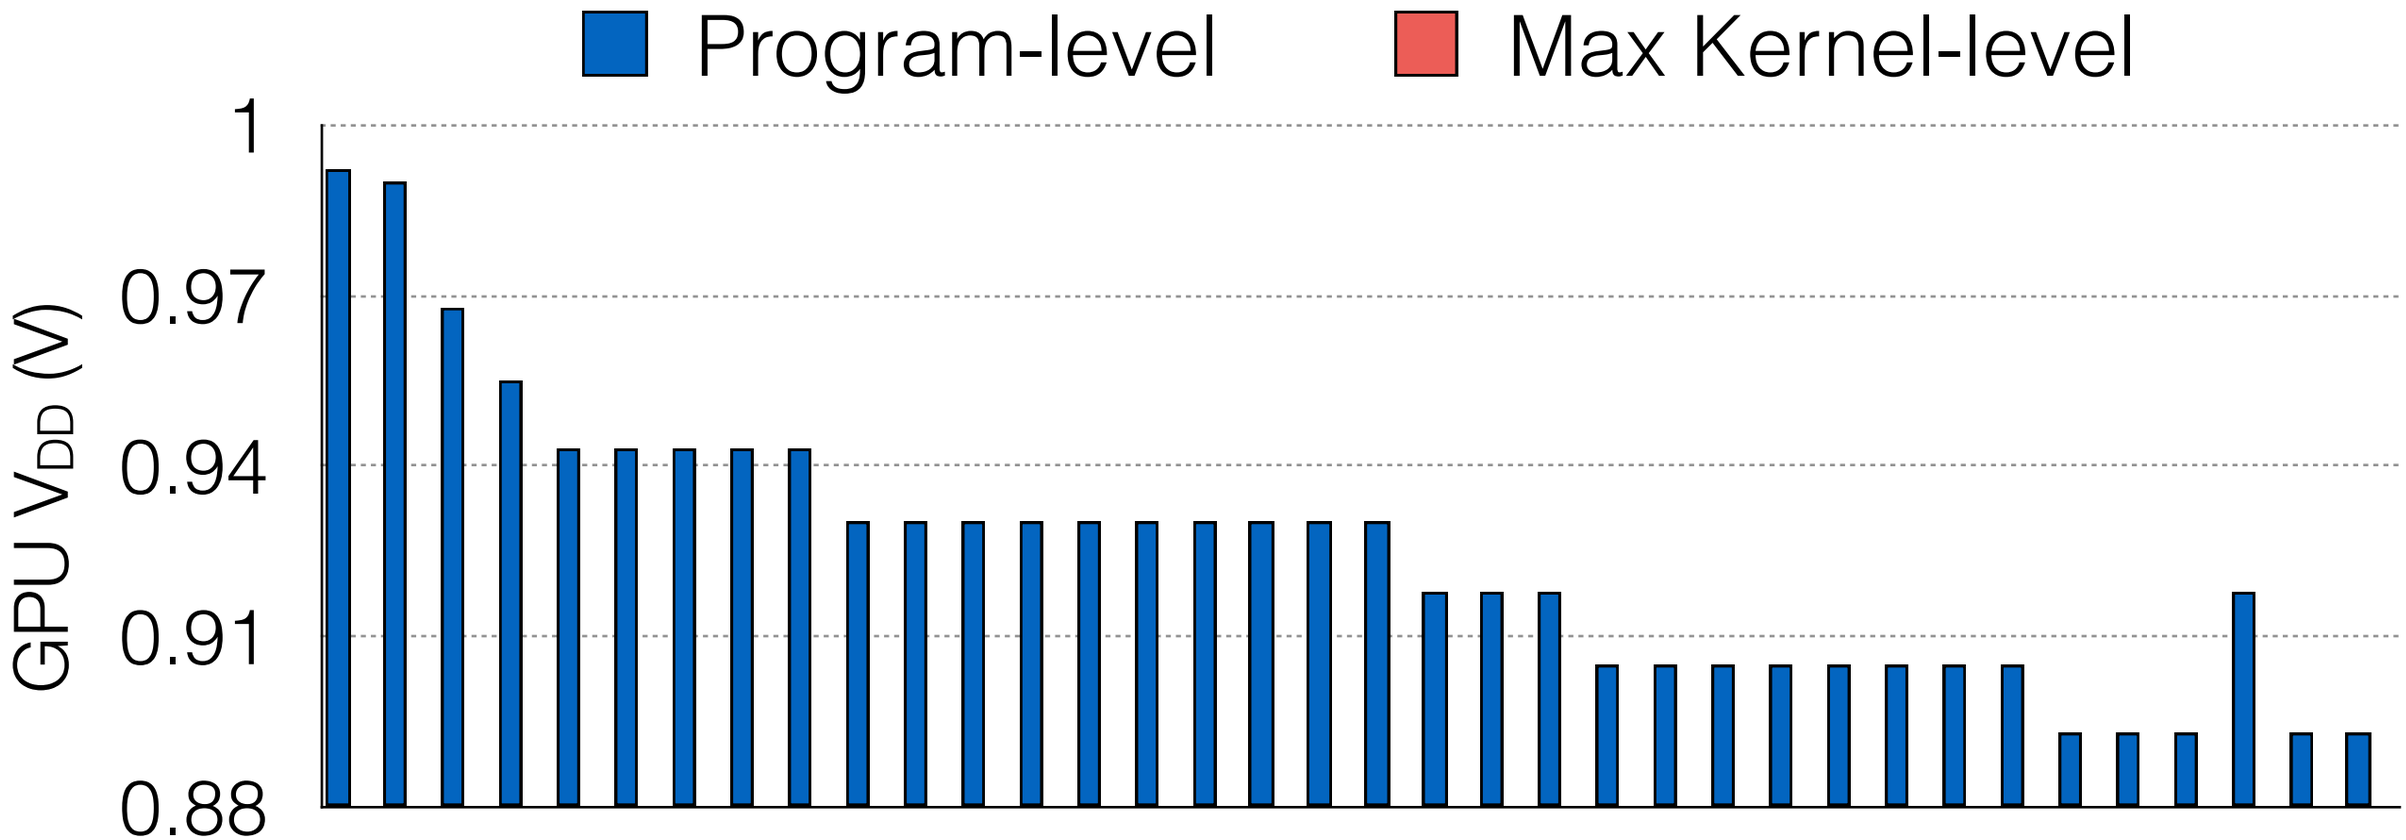

Program/Kernel Level V_{min} Comparison

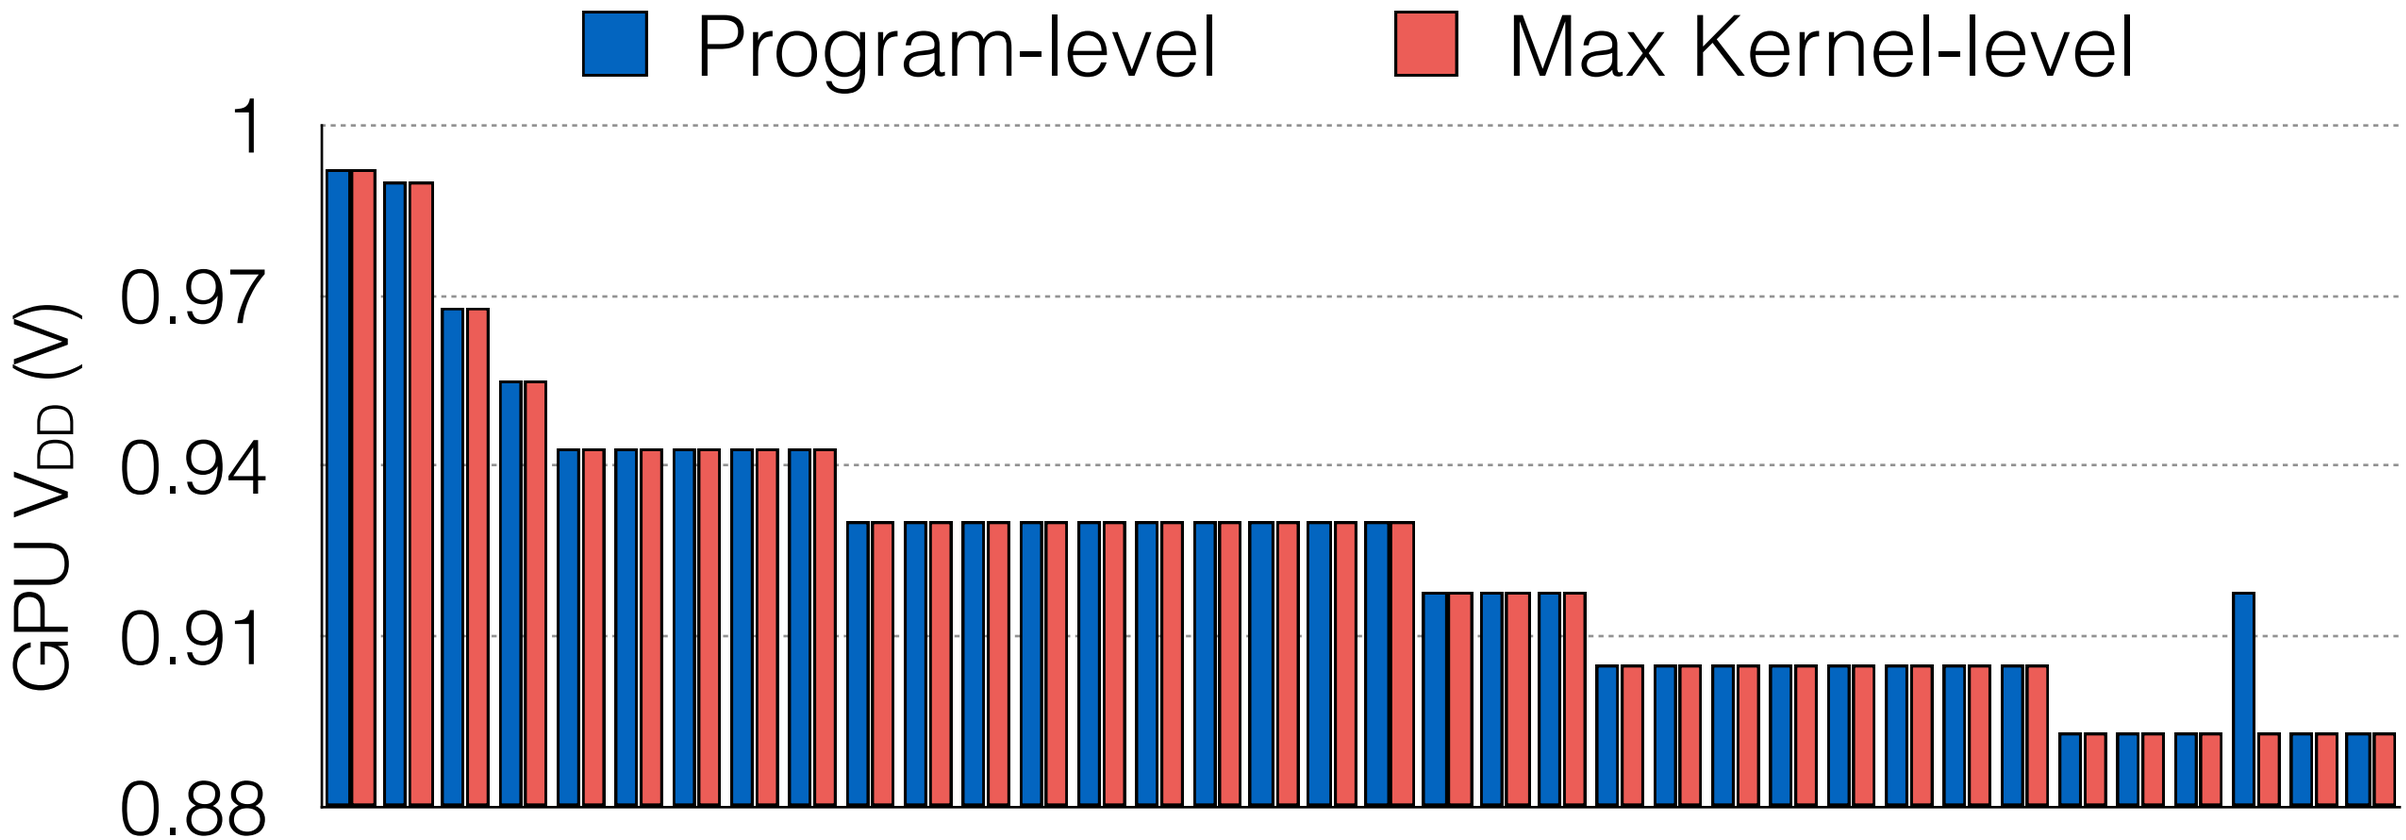

Program/Kernel Level V_{min} Comparison

- Program-level V_{min} same as maximum kernel-level V_{min}

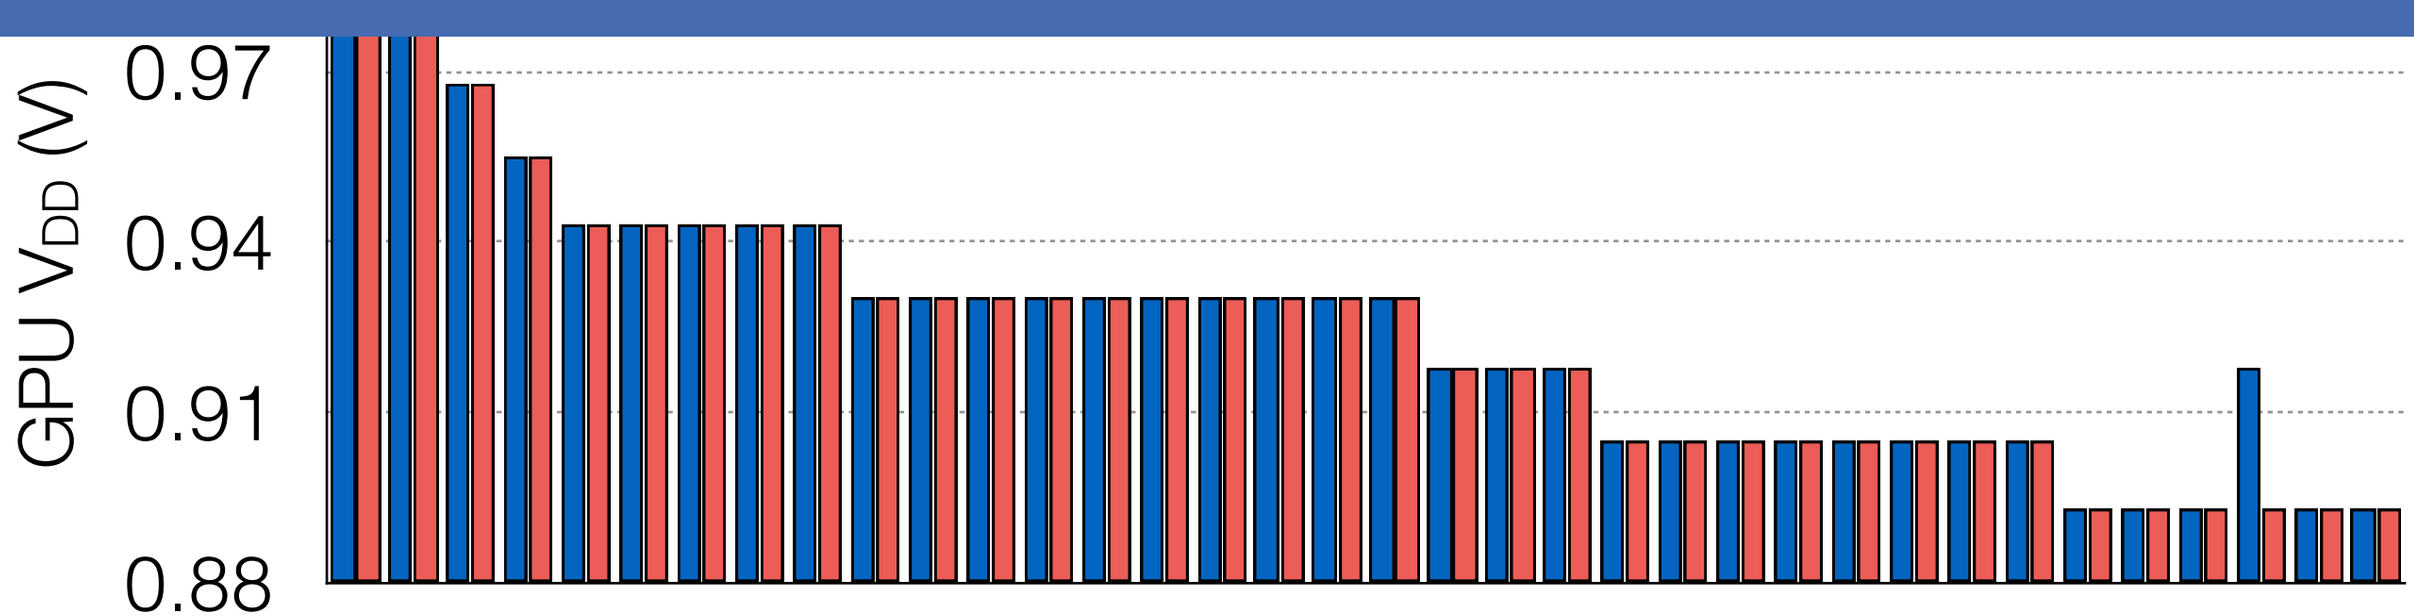

Program/Kernel Level V_{min} Comparison

- Program-level V_{min} same as maximum kernel-level V_{min}
- Inter-kernel activity does not determine V_{min} value

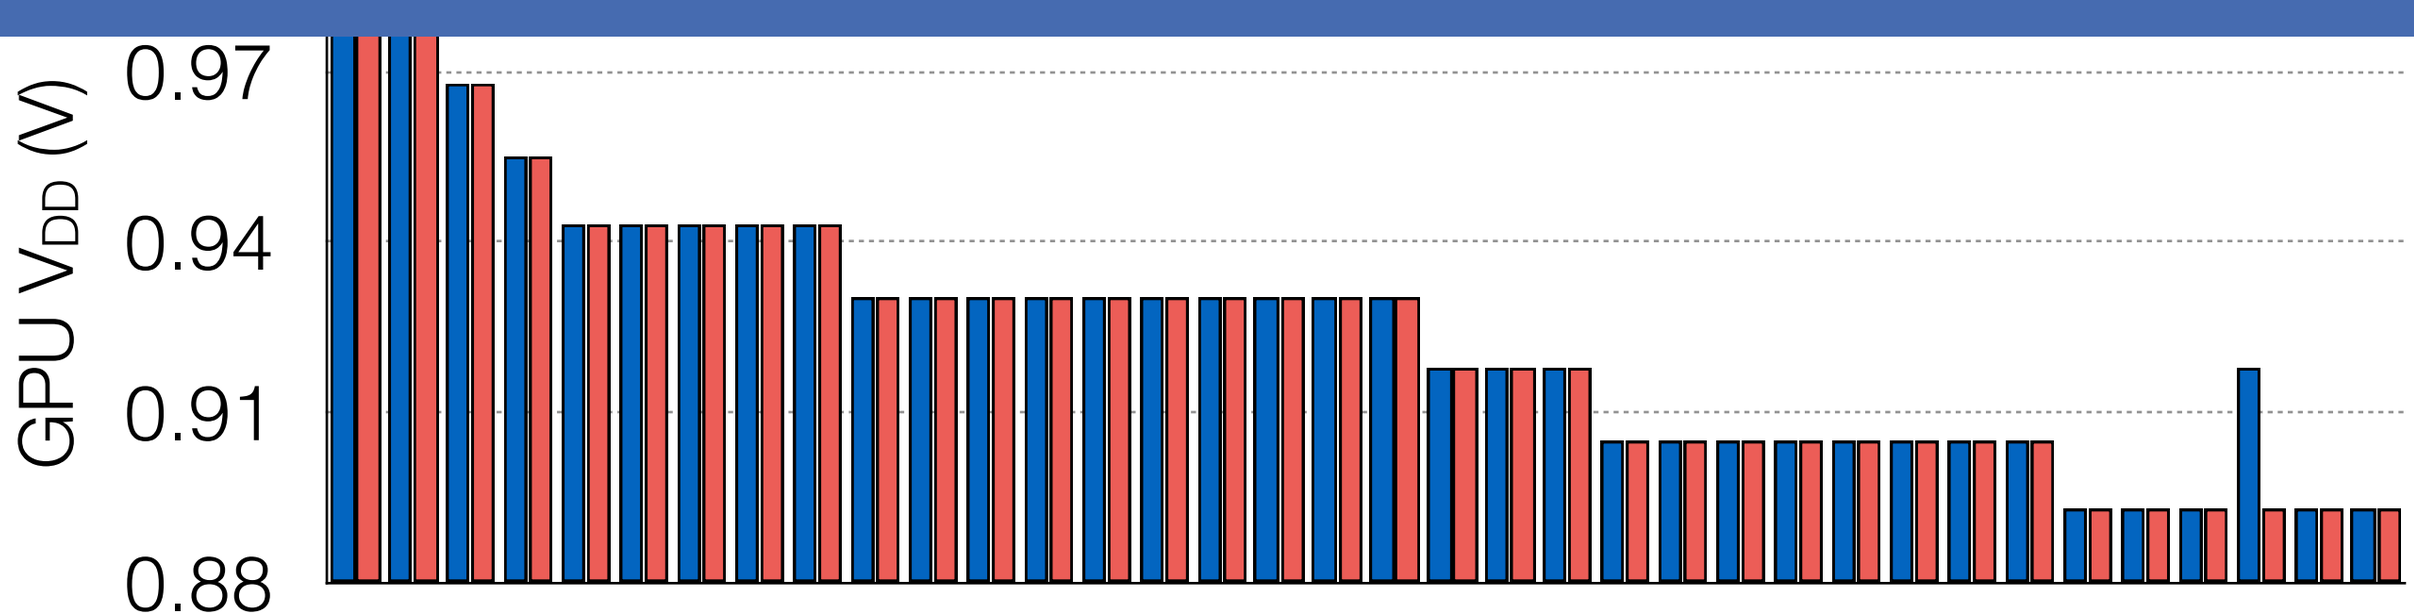

Program/Kernel Level V_{min} Comparison

- Program-level V_{min} same as maximum kernel-level V_{min}
- Inter-kernel activity does not determine V_{min} value

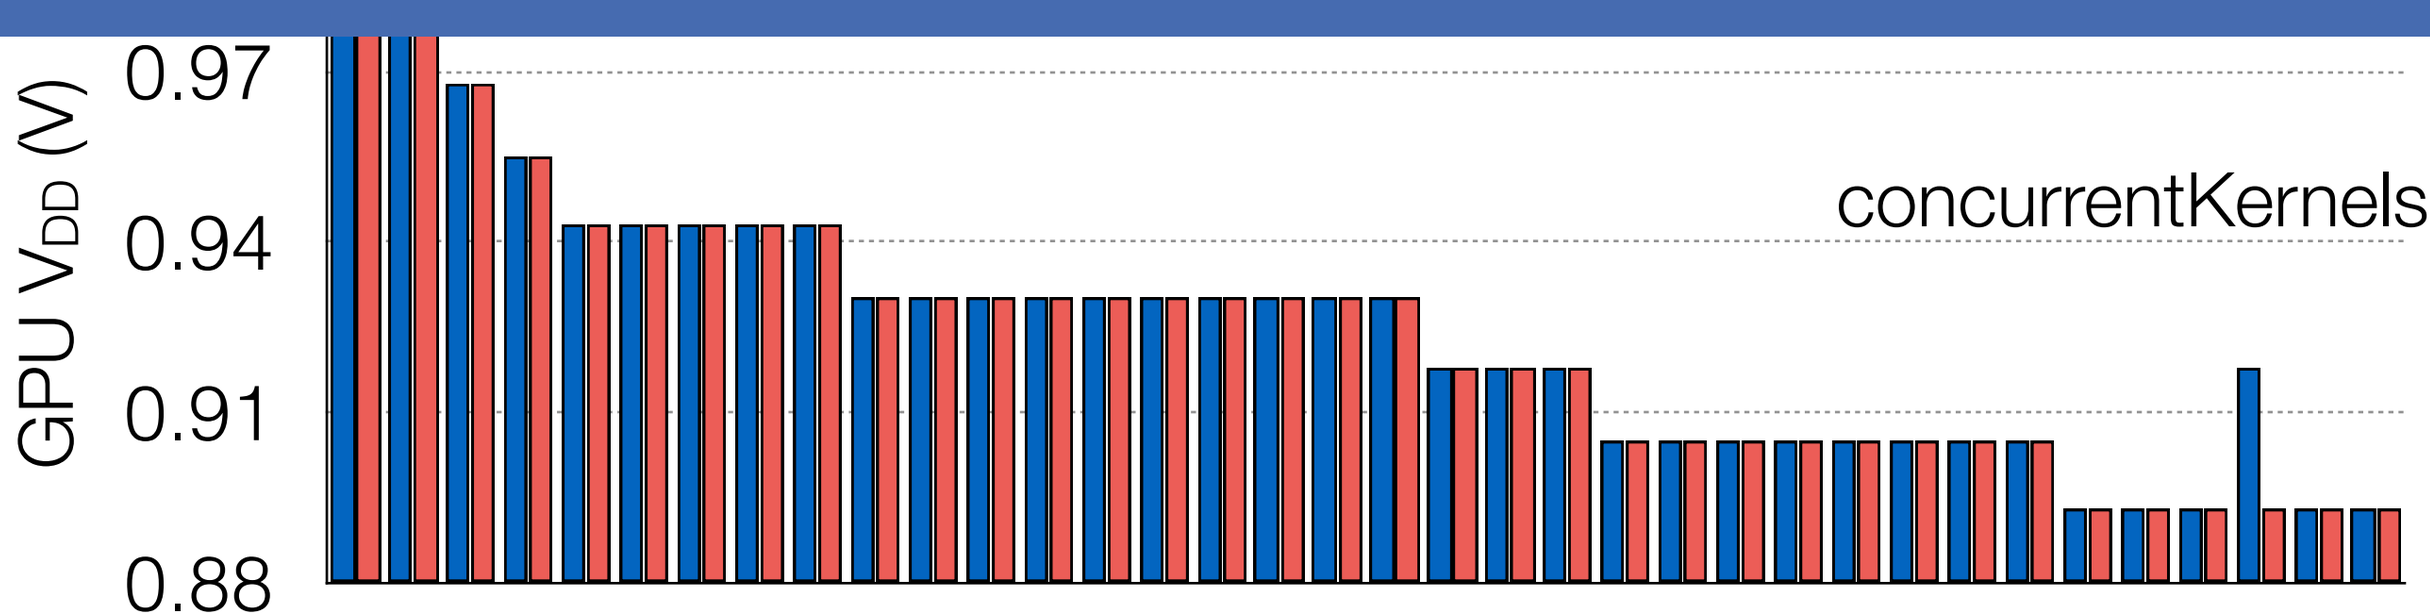

Program/Kernel Level V_{min} Comparison

- Program-level V_{min} same as maximum kernel-level V_{min}
- Inter-kernel activity does not determine V_{min} value

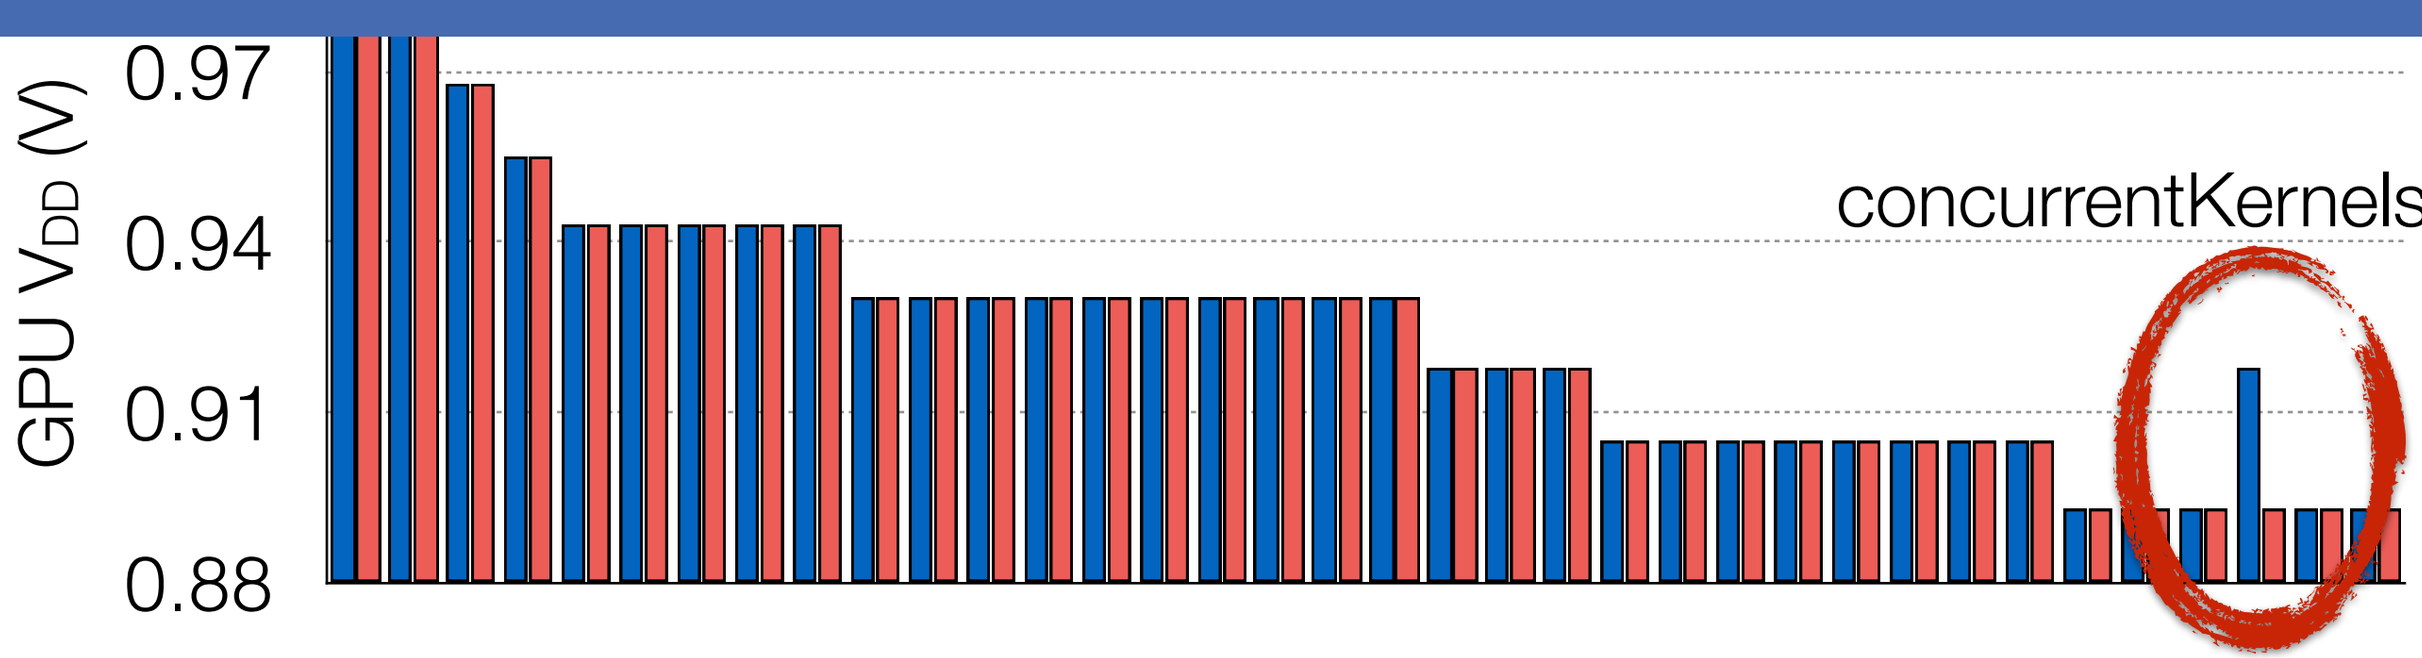

Executive Summary

- Guardband measurement
- Guardband analysis
- Guardband optimization

Predictive Guardbanding

Predictive Guardbanding

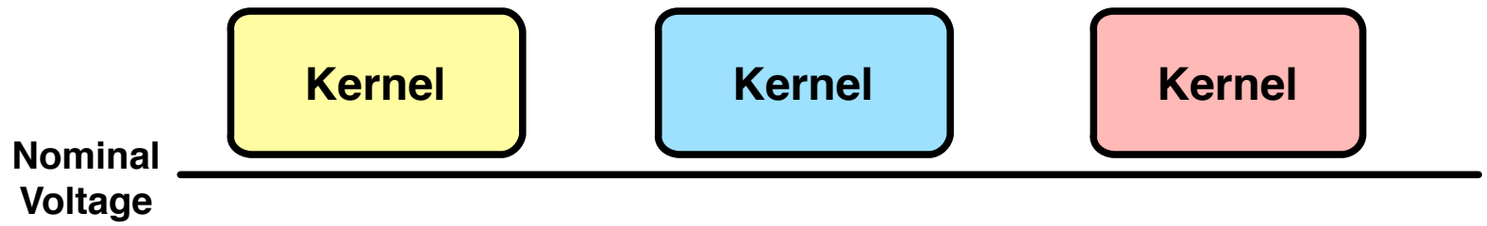

Predictive Guardbanding

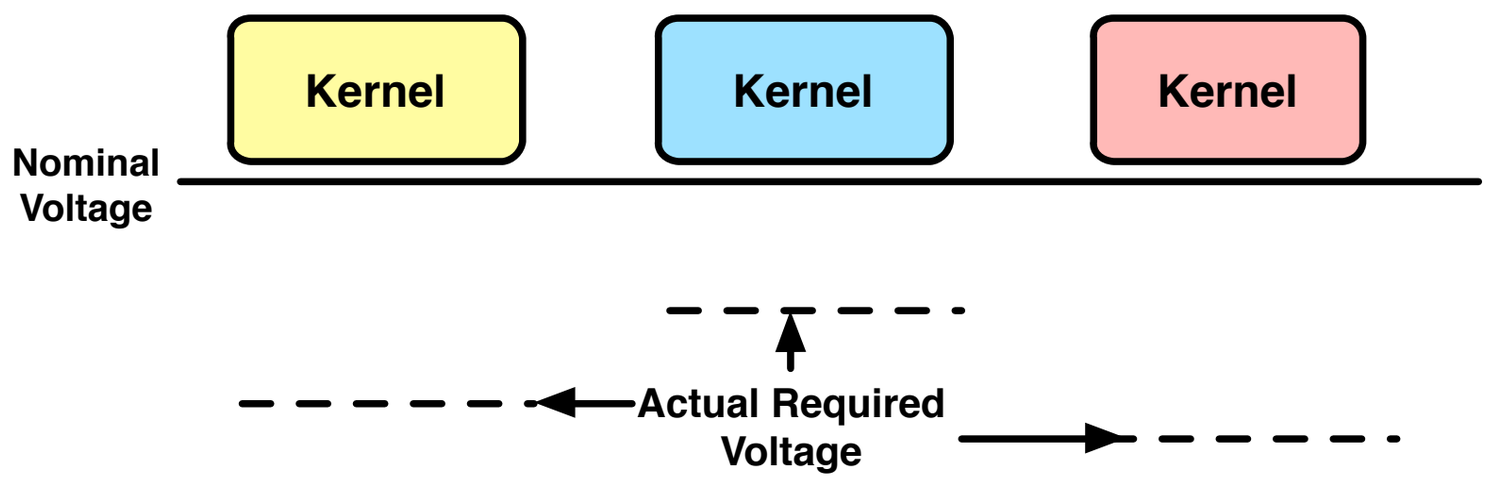

Predictive Guardbanding

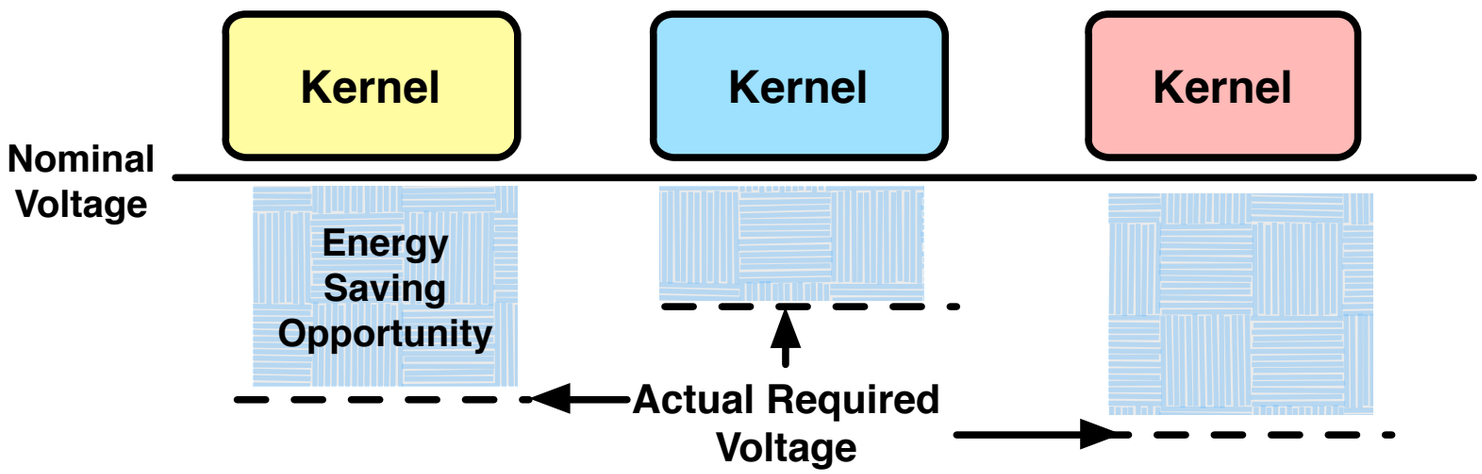

Predictive Guardbanding

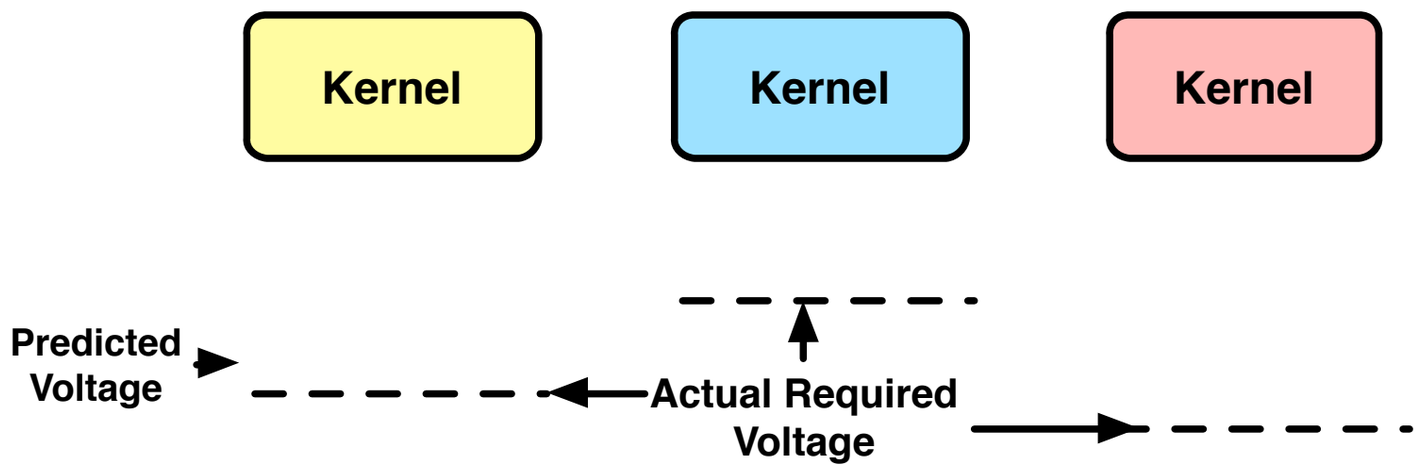

Predictive Guardbanding

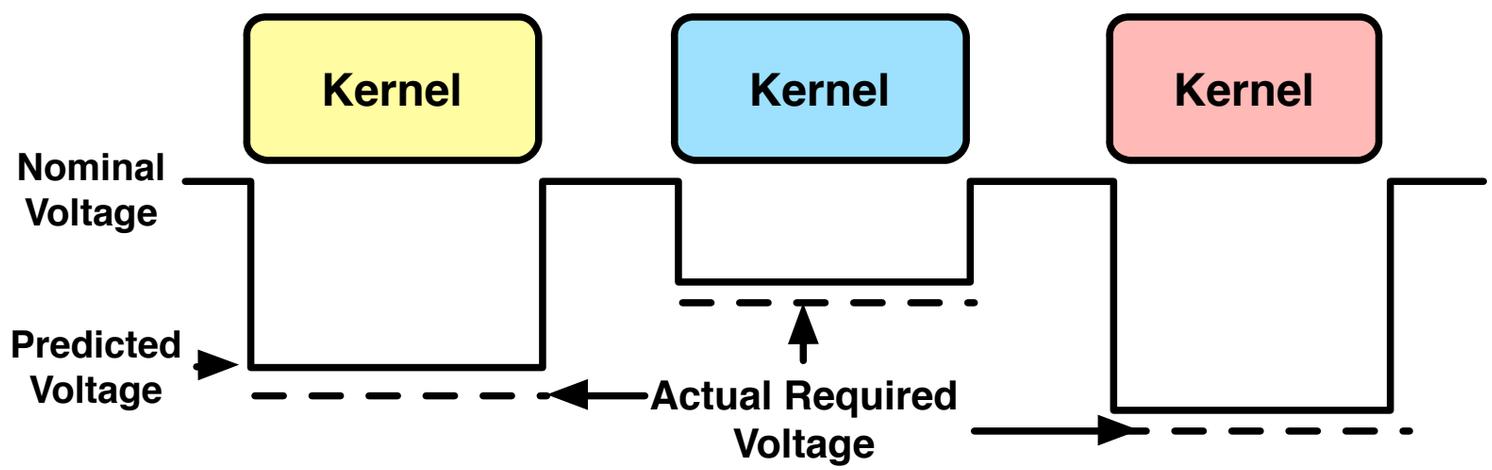

Predictive Guardbanding

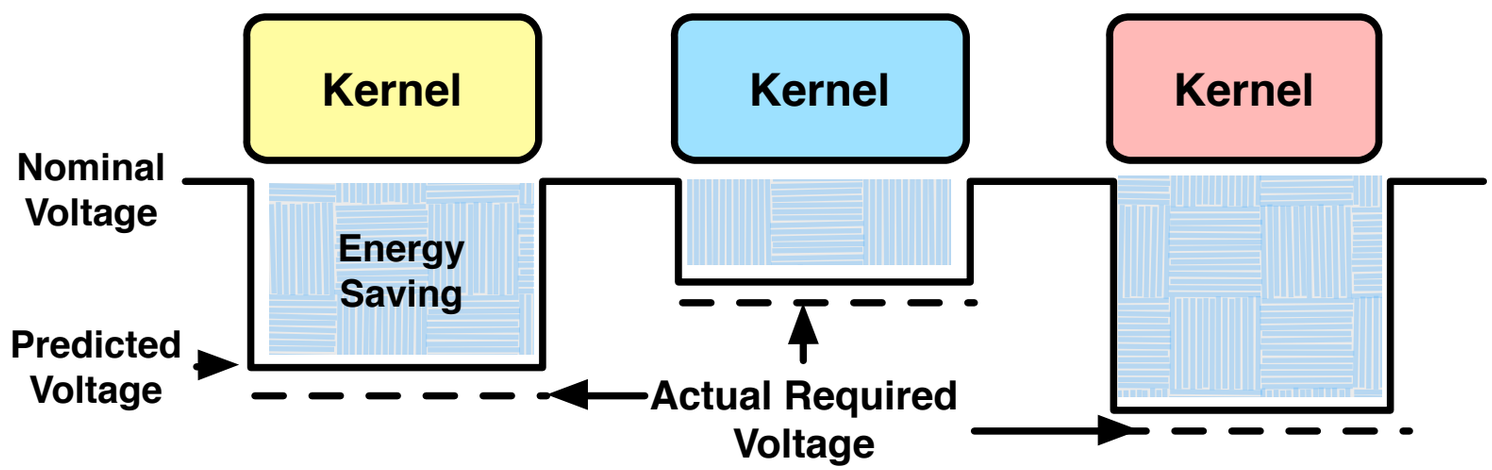

Predictive Guardbanding

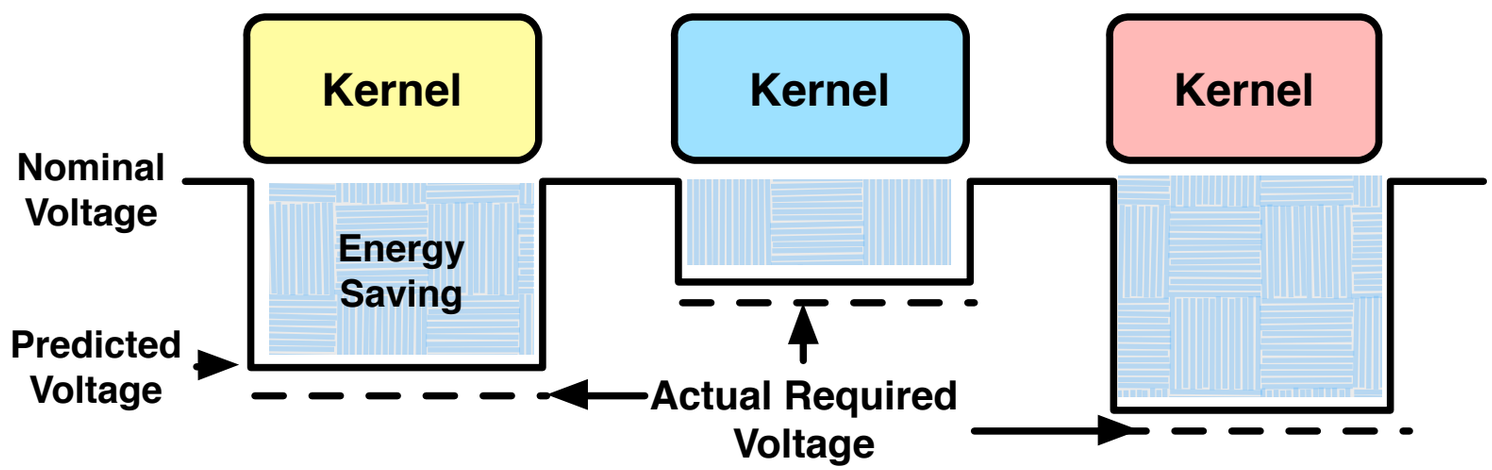

- Exploit program-dependent V_{min} behavior
- Program/kernel level V_{min} prediction

Performance Counter Based V_{min} Prediction

Performance Counter Based V_{min} Prediction

- Use all available performance counters to construct a V_{min} prediction model

Performance Counter Based V_{min} Prediction

- Use all available performance counters to construct a V_{min} prediction model

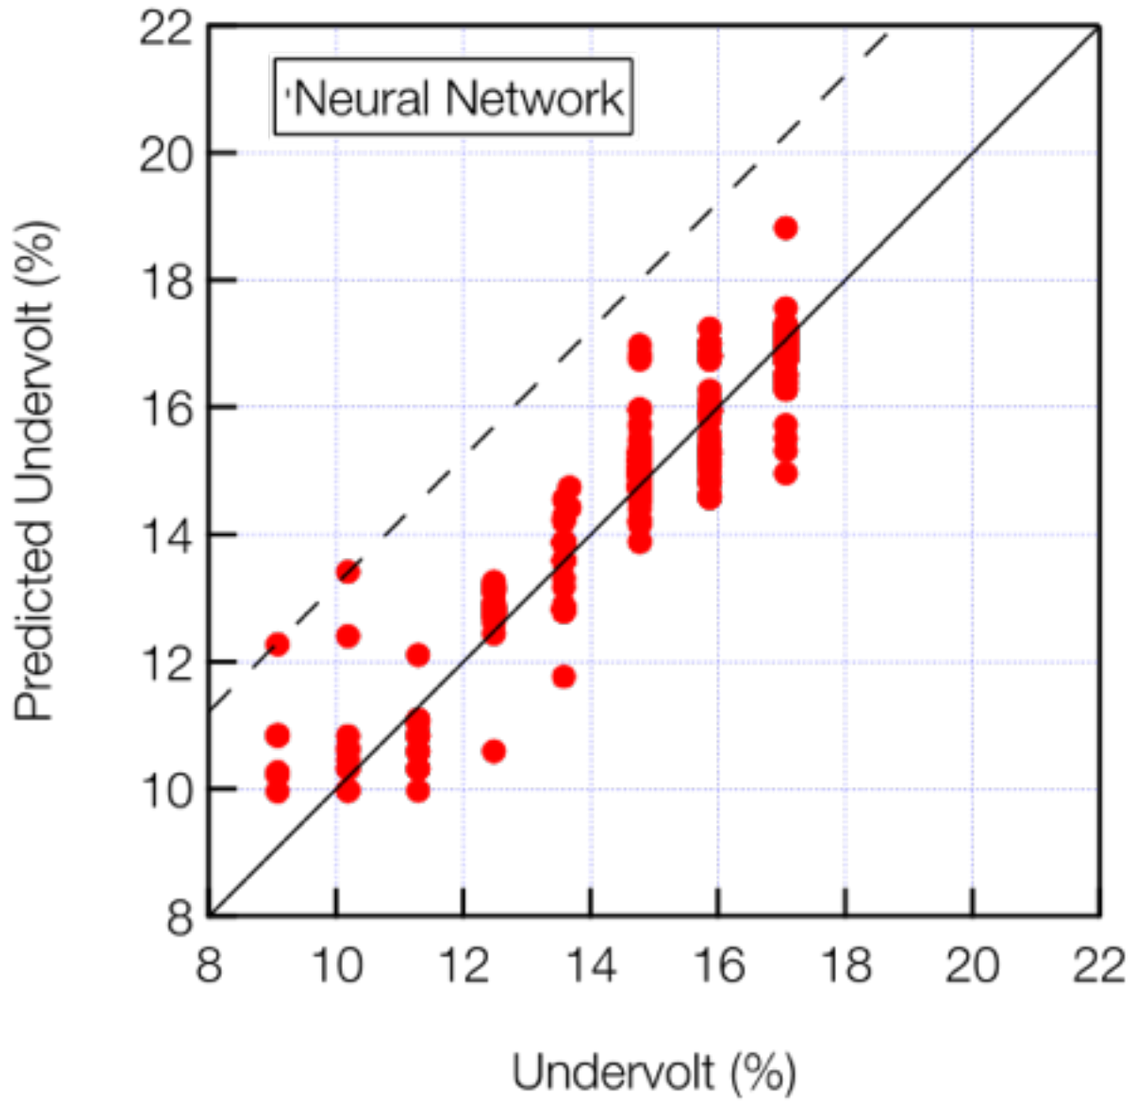

Neural network
RMSE: 0.5%, max error: 3%

Energy Efficiency Optimization Potential

Energy Efficiency Optimization Potential

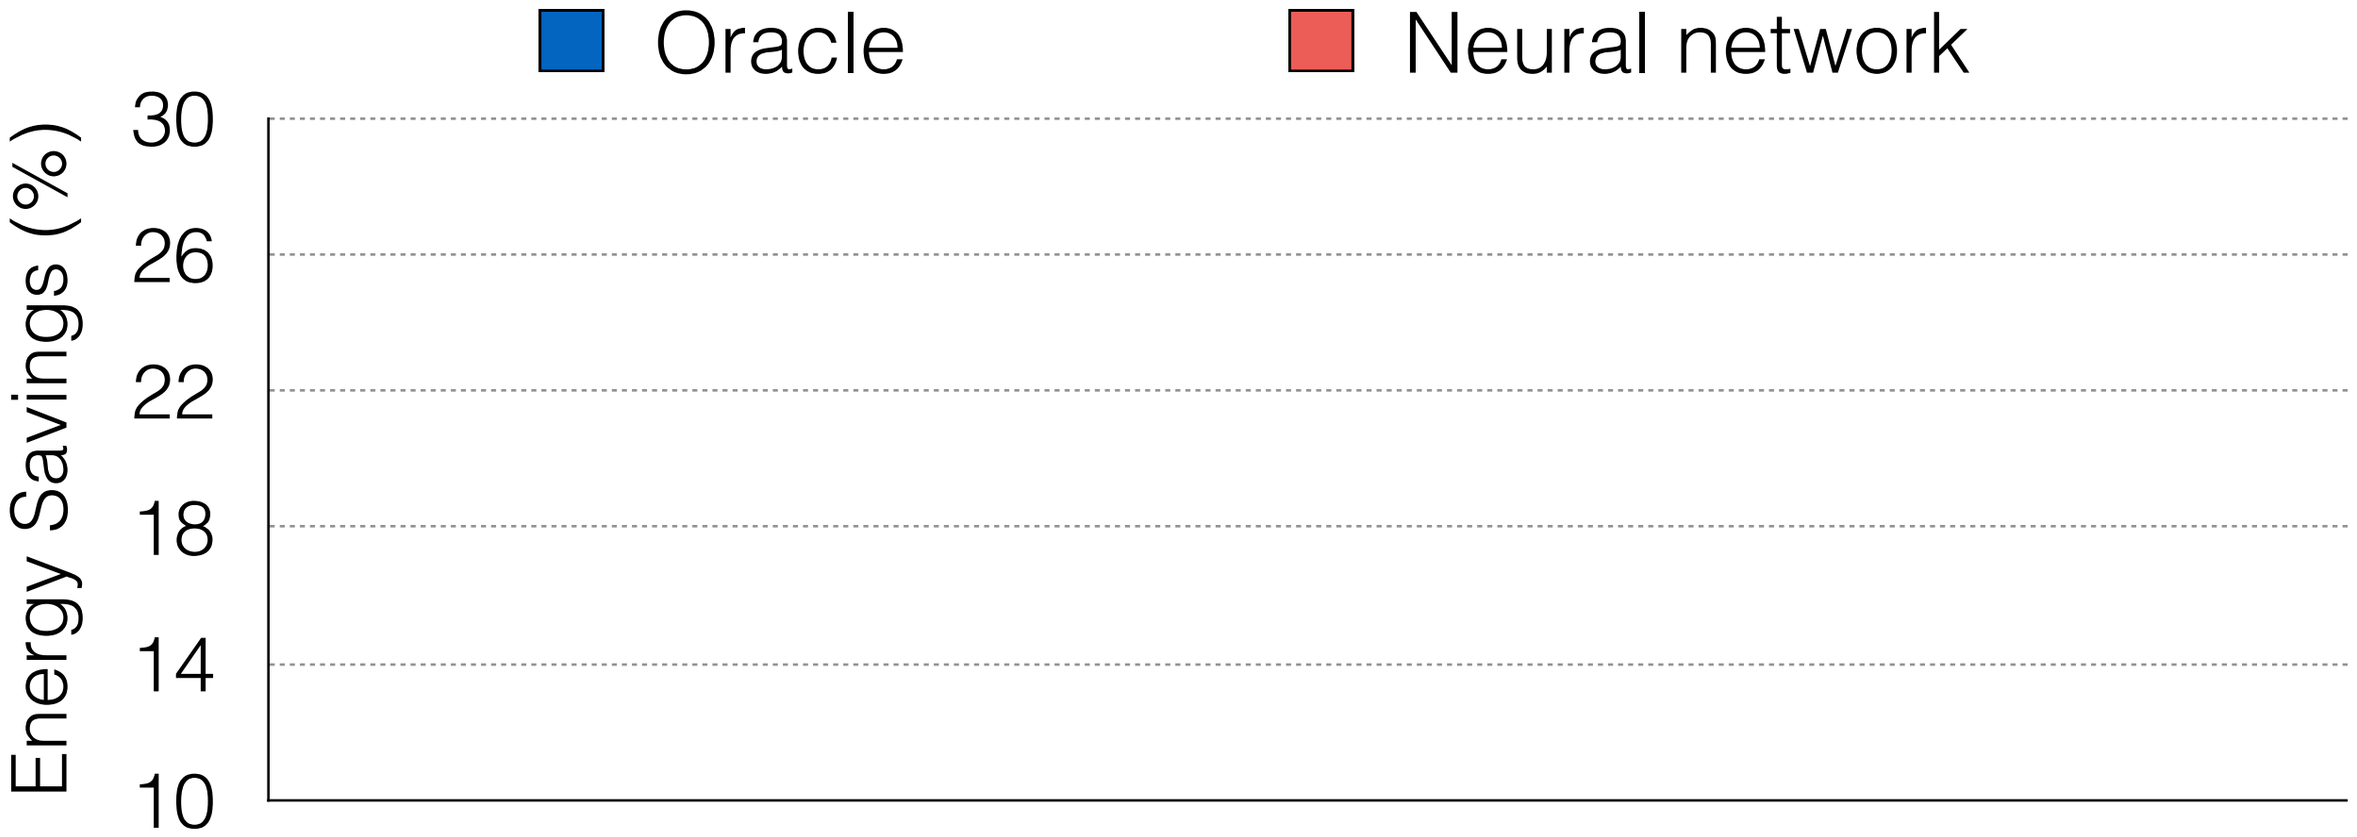

Energy Efficiency Optimization Potential

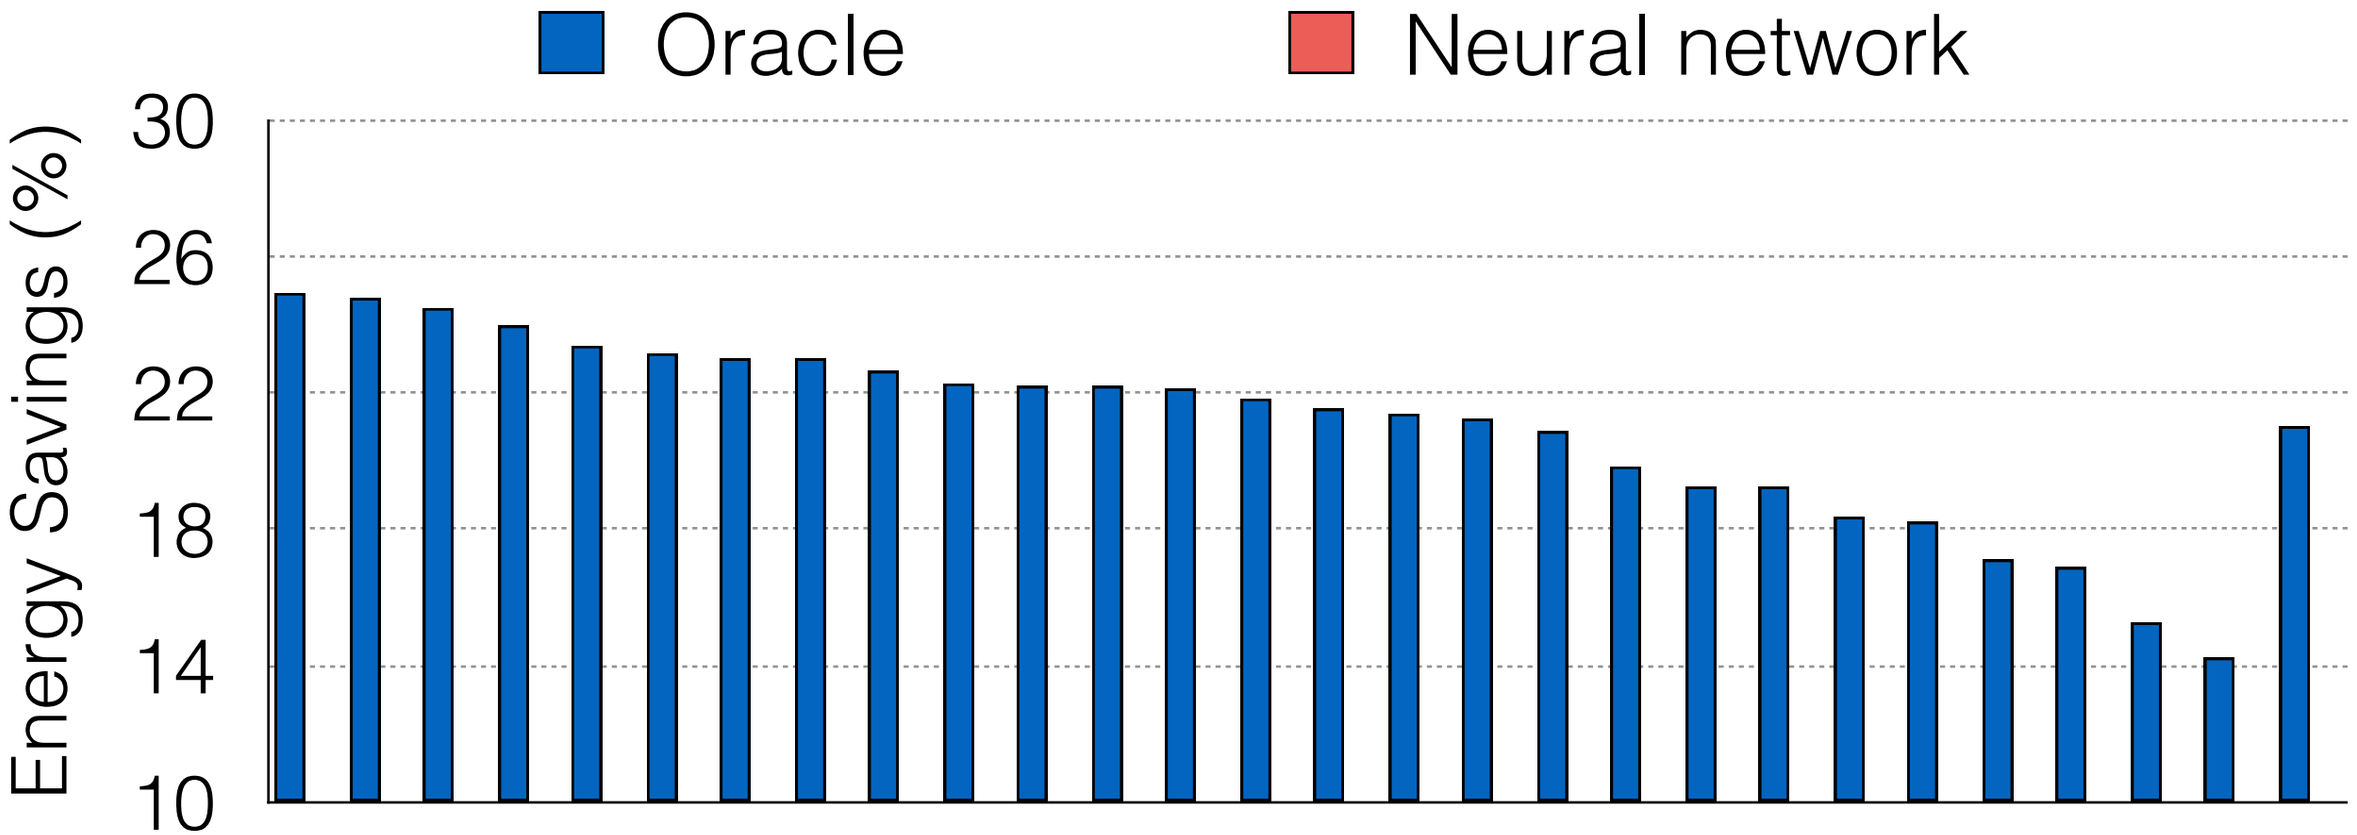

Energy Efficiency Optimization Potential

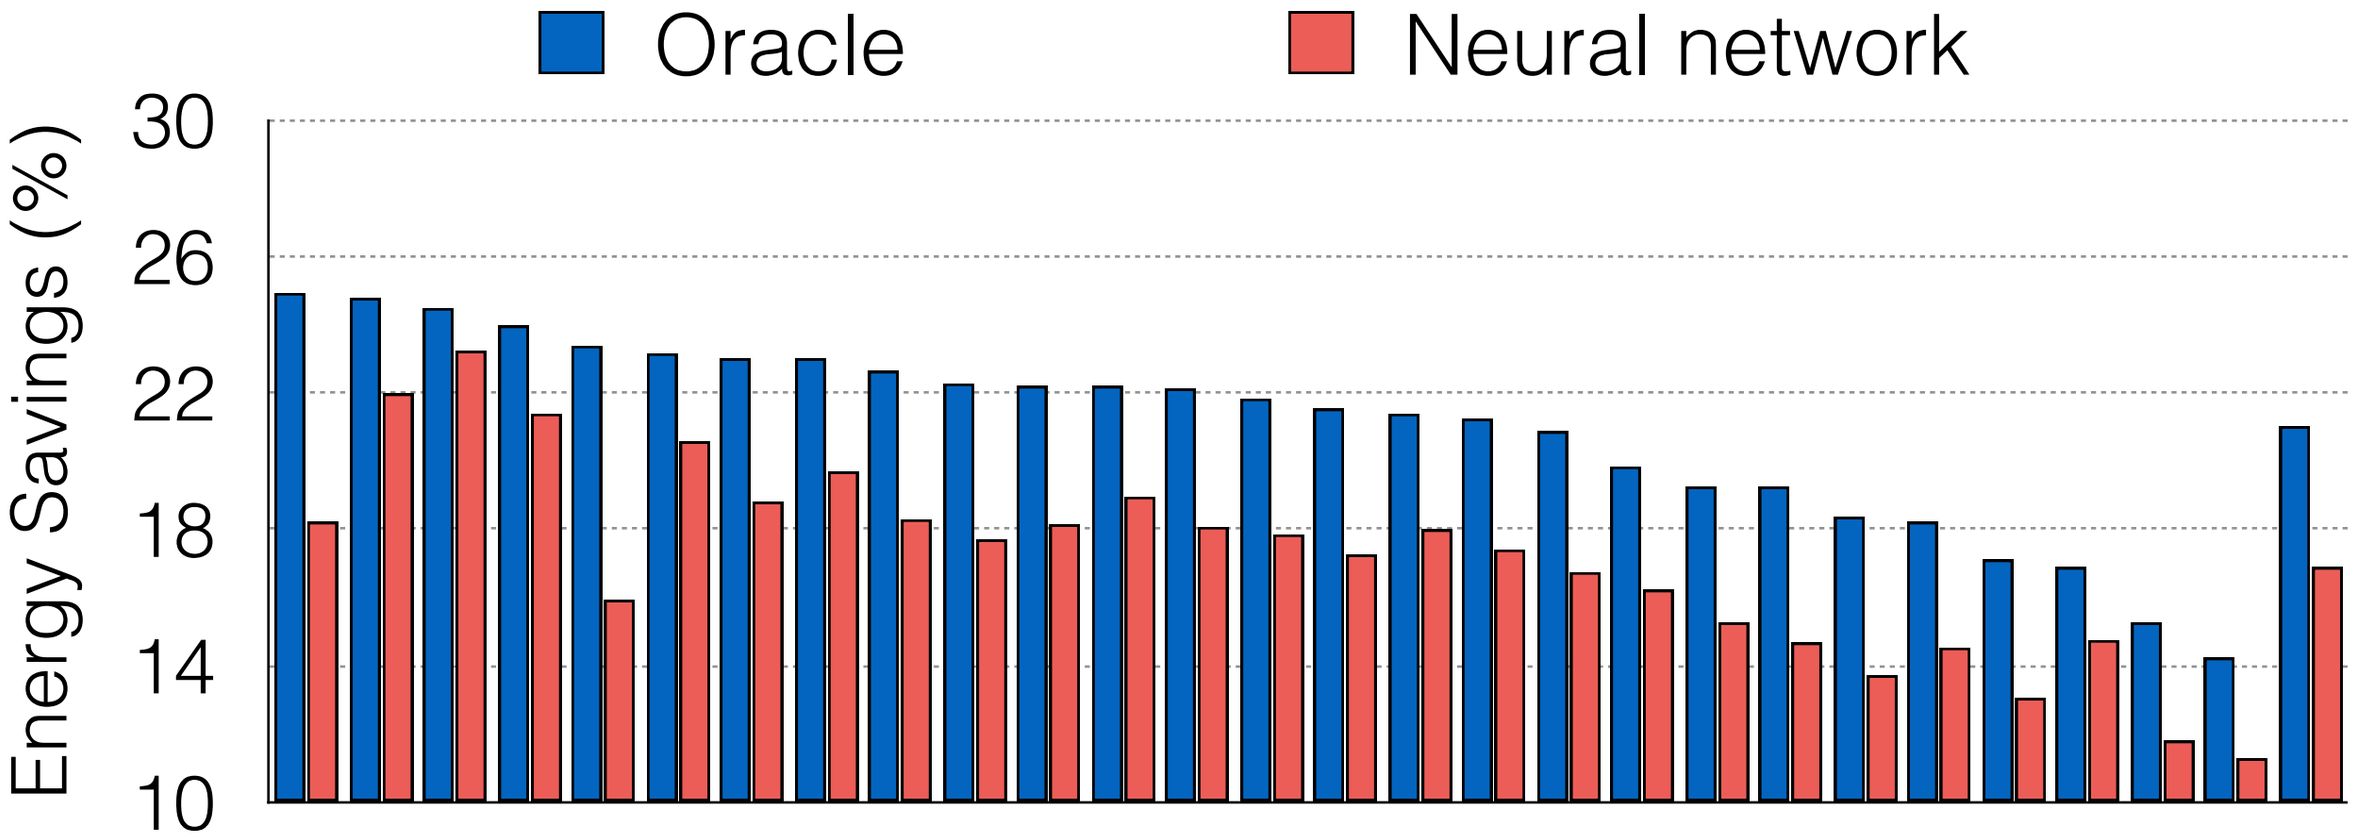

Energy Efficiency Optimization Potential

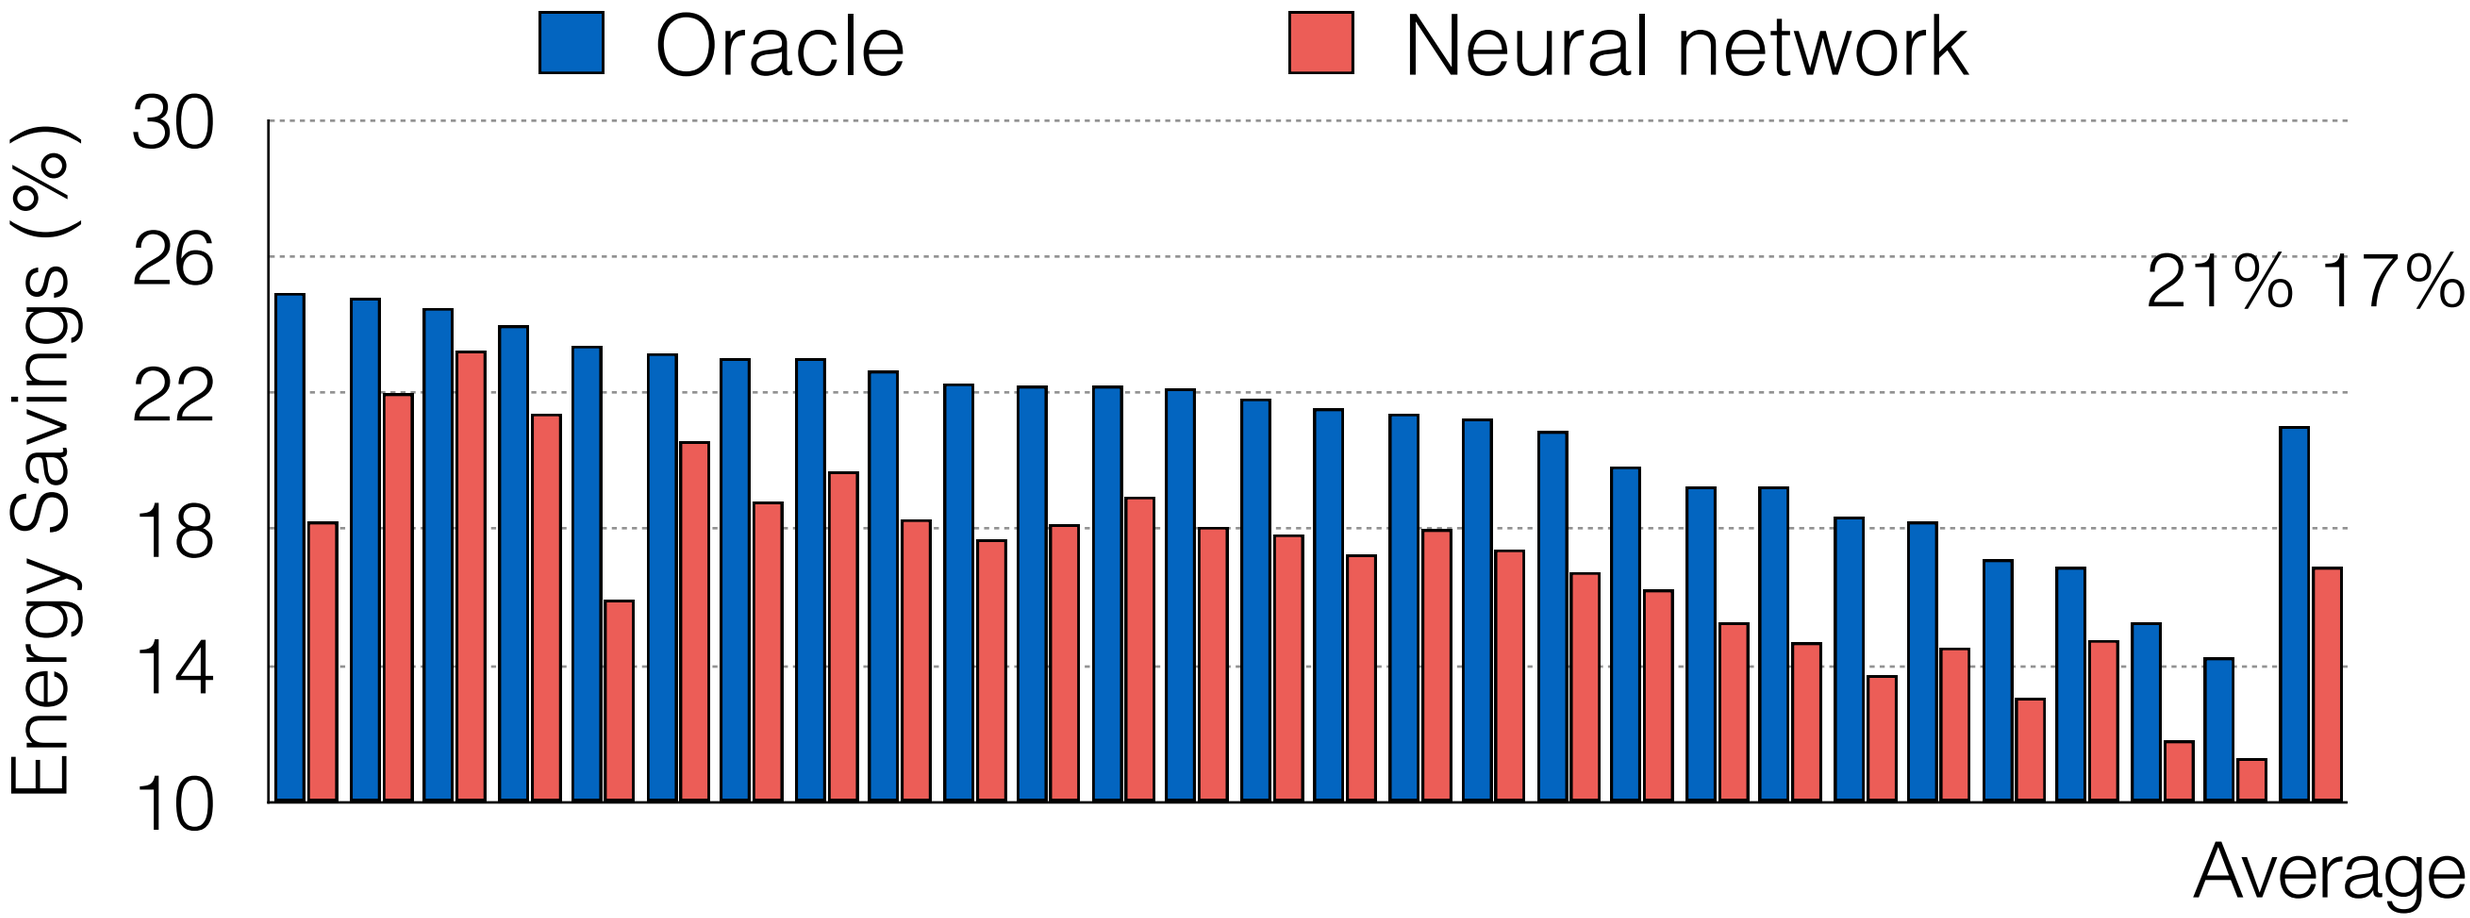

Energy Efficiency Optimization Potential

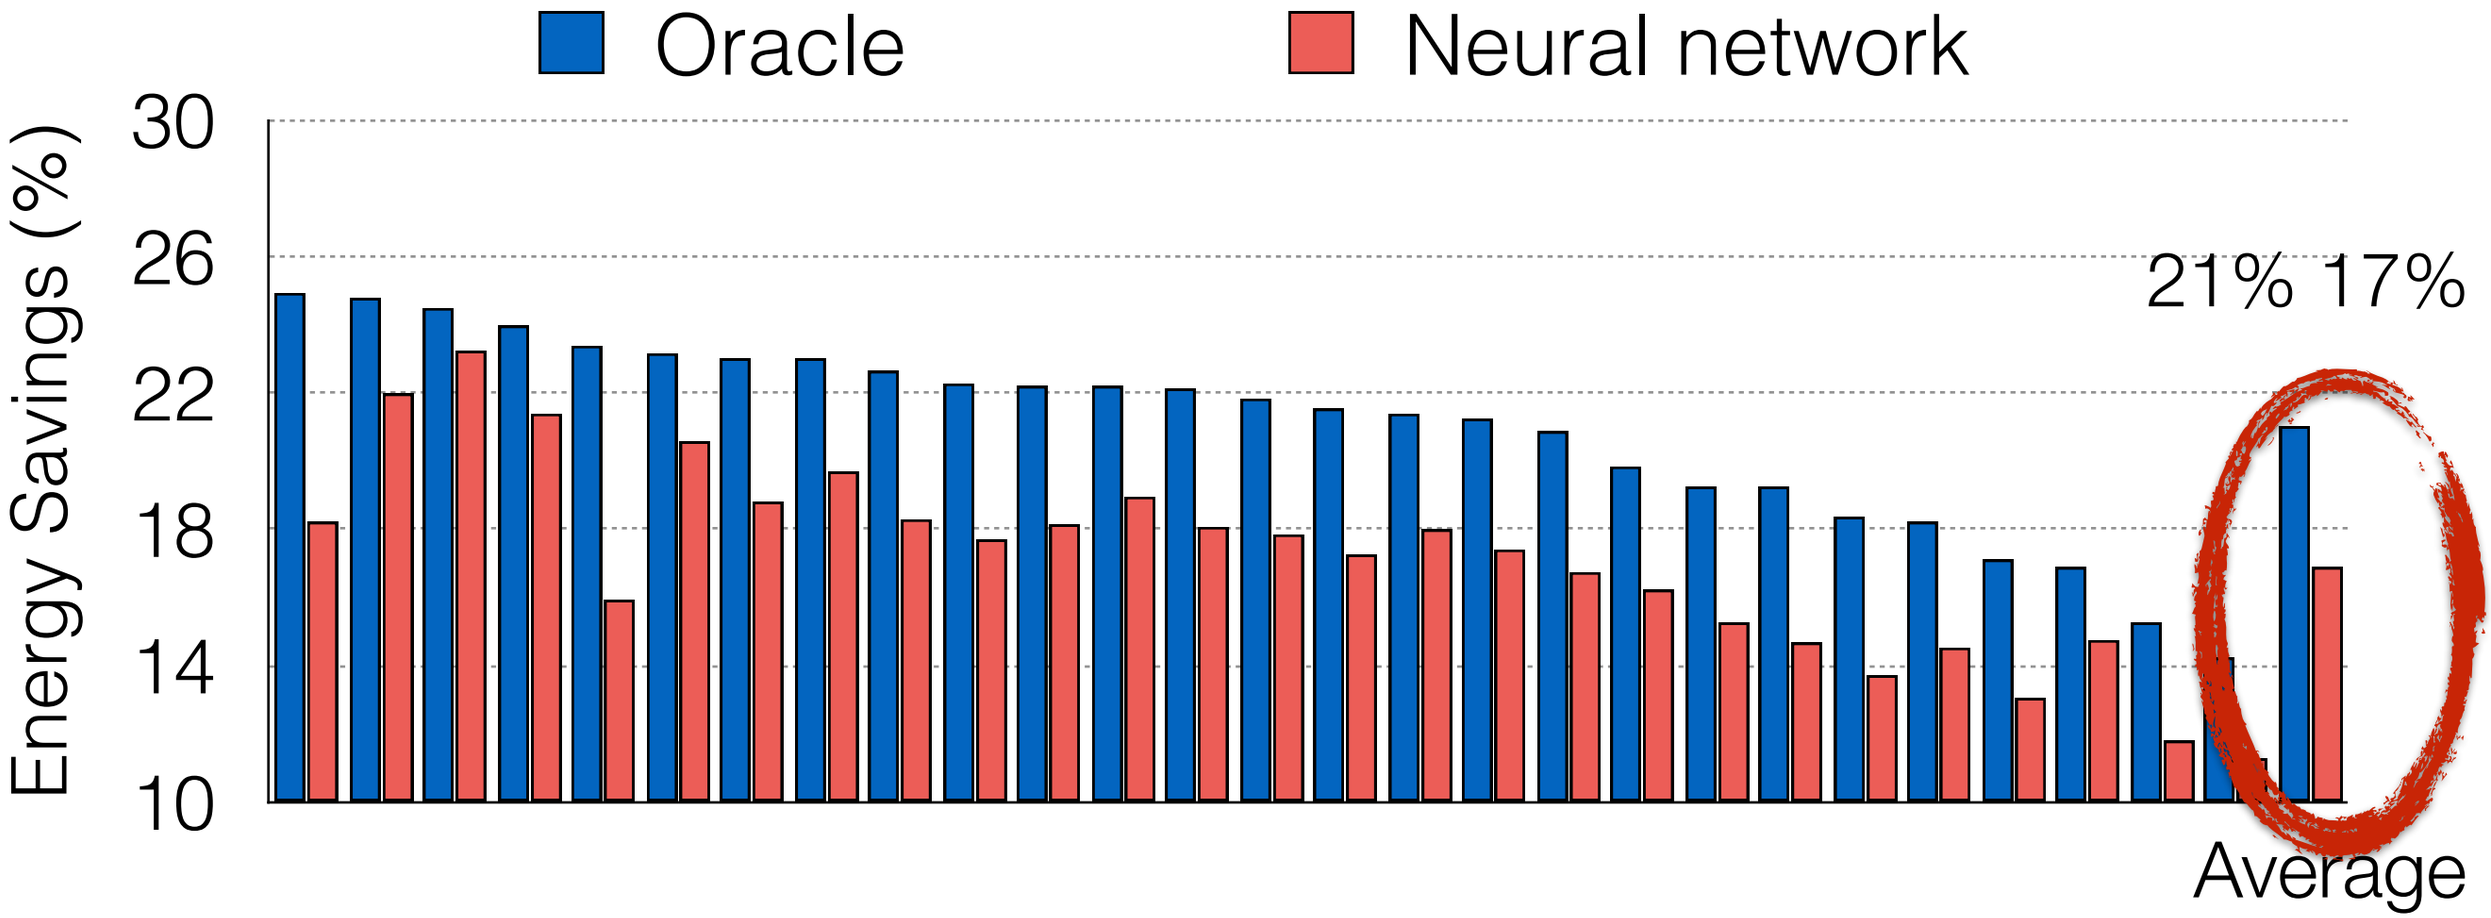

Conclusion

Conclusion

Large amount (up to 20%) of voltage guardband for GPUs

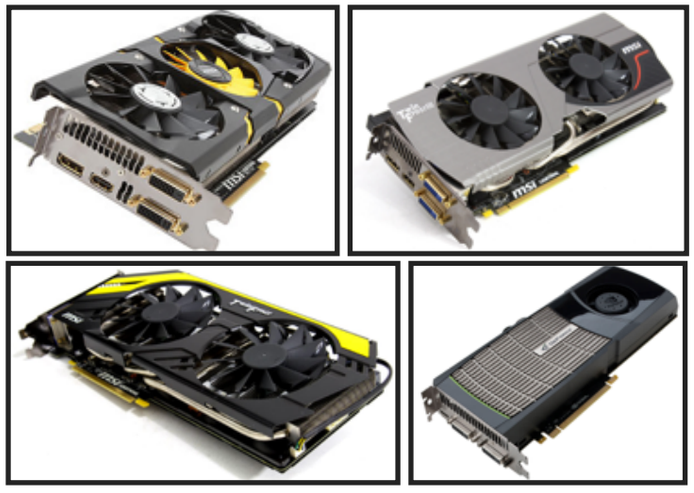

Conclusion

Large amount (up to 20%) of voltage guardband for GPUs

Intra-kernel di/dt droop is the largest guardband determinant

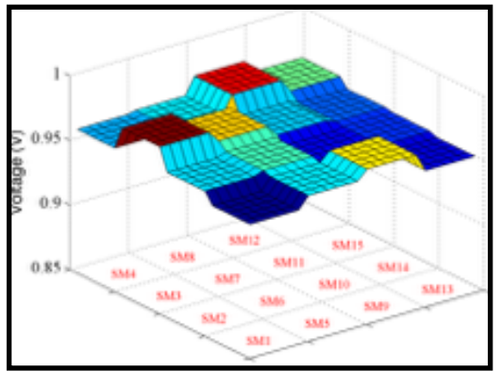

Conclusion

Large amount (up to 20%) of voltage guardband for GPUs

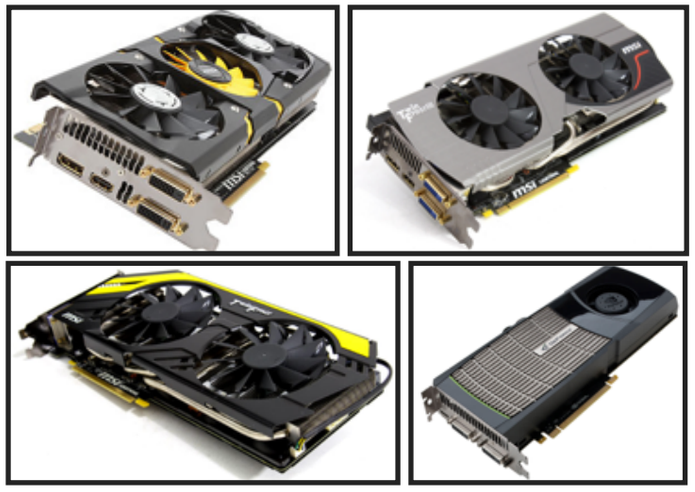

Intra-kernel di/dt droop is the largest guardband determinant

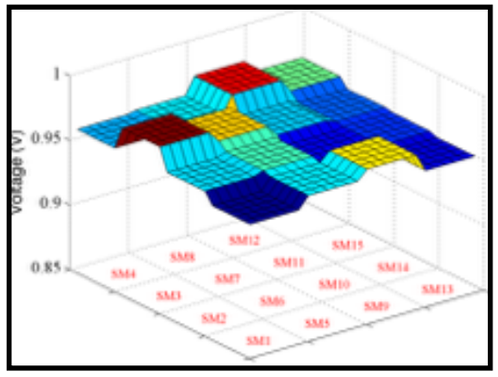

We show the potential of program-driven predictive guardbanding