www.phoronix-test-suite.com

pytorch hx

AMD Ryzen 9 5900HX testing with a ASUS G513QY v1.0 (G513QY.318 BIOS) and ASUS AMD Cezanne 512MB on Ubuntu 22.10 via the Phoronix Test Suite.

Automated Executive Summary

b had the most wins, coming in first place for 44% of the tests.

Based on the geometric mean of all complete results, the fastest (b) was 1.028x the speed of the slowest (c). a was 0.982x the speed of b and c was 0.99x the speed of a.

Test Systems:

a

b

pytorch hx

c

Processor: AMD Ryzen 9 5900HX @ 3.30GHz (8 Cores / 16 Threads), Motherboard: ASUS G513QY v1.0 (G513QY.318 BIOS), Chipset: AMD Renoir/Cezanne, Memory: 16GB, Disk: 512GB SAMSUNG MZVLQ512HBLU-00B00, Graphics: ASUS AMD Cezanne 512MB (2500/1000MHz), Audio: AMD Navi 21/23, Monitor: LQ156M1JW25, Network: Realtek RTL8111/8168/8411 + MEDIATEK MT7921 802.11ax PCI

OS: Ubuntu 22.10, Kernel: 5.19.0-46-generic (x86_64), Desktop: GNOME Shell 43.0, Display Server: X Server 1.21.1.4 + Wayland, OpenGL: 4.6 Mesa 22.2.5 (LLVM 15.0.2 DRM 3.47), Vulkan: 1.3.224, Compiler: GCC 12.2.0, File-System: ext4, Screen Resolution: 1920x1080

Kernel Notes: Transparent Huge Pages: madvise
Processor Notes: Scaling Governor: acpi-cpufreq schedutil (Boost: Enabled) - Platform Profile: balanced - CPU Microcode: 0xa50000c - ACPI Profile: balanced
Python Notes: Python 3.10.7
Security Notes: itlb_multihit: Not affected + l1tf: Not affected + mds: Not affected + meltdown: Not affected + mmio_stale_data: Not affected + retbleed: Not affected + spec_store_bypass: Mitigation of SSB disabled via prctl + spectre_v1: Mitigation of usercopy/swapgs barriers and __user pointer sanitization + spectre_v2: Mitigation of Retpolines IBPB: conditional IBRS_FW STIBP: always-on RSB filling PBRSB-eIBRS: Not affected + srbds: Not affected + tsx_async_abort: Not affected

	a	b	c
PyTorch - CPU - 1 - ResNet-50 (batches/sec)	34.83	**38.17**	**33.69**
Normalized	91.25%	100%	88.26%
Standard Deviation	1.2%		
PyTorch - CPU - 512 - ResNet-50	18.38	**19.20**	**17.58**
Normalized	95.73%	100%	91.56%
Standard Deviation	1.1%		
PyTorch - CPU - 16 - ResNet-152	**8.99**	9.07	**9.42**
Normalized	95.44%	96.28%	100%
Standard Deviation	2.4%		
PyTorch - CPU - 16 - ResNet-50	**19.54**	19.41	**18.72**
Normalized	100%	99.33%	95.8%
Standard Deviation	3.3%		
PyTorch - CPU - 512 - ResNet-152 (batches/sec)	**8.74**	8.72	**8.50**
Normalized	100%	99.77%	97.25%
Standard Deviation	0.8%		
PyTorch - CPU - 1 - ResNet-152	**14.90**	**15.25**	15.22
Normalized	97.7%	100%	99.8%
Standard Deviation	2.1%		
PyTorch - CPU - 512 - Efficientnet_v2_l (batches/sec)	6.23	**6.29**	**6.17**
Normalized	99.05%	100%	98.09%
Standard Deviation	1.4%		
PyTorch - CPU - 1 - Efficientnet_v2_l (batches/sec)	9.37	**9.36**	**9.41**
Normalized	99.57%	99.47%	100%
Standard Deviation	1%		
PyTorch - CPU - 16 - Efficientnet_v2_l (batches/sec)	**6.28**	**6.27**	**6.28**
Normalized	100%	99.84%	100%
Standard Deviation	0.8%		

pytorch hx

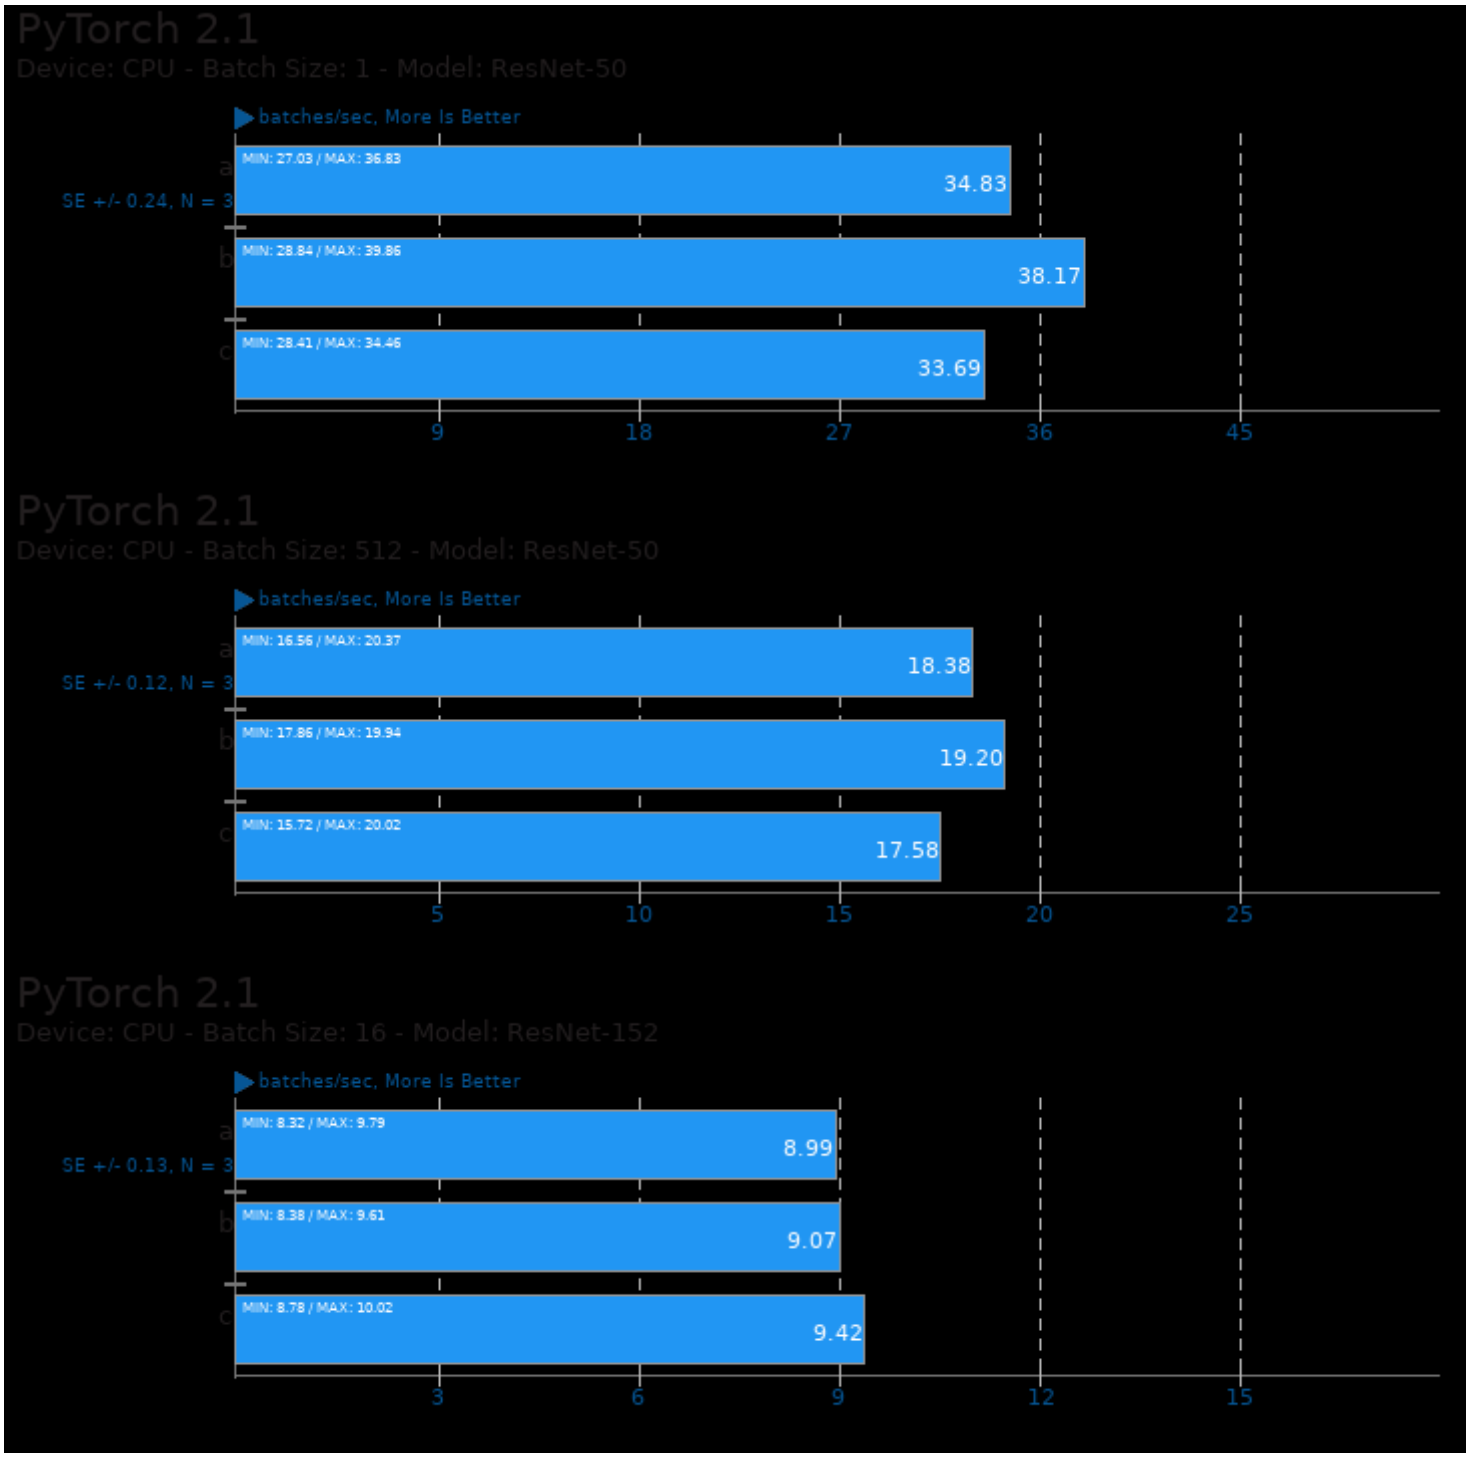
PyTorch 2.1
Device: CPU - Batch Size: 1 - Model: ResNet-50
batches/sec, More Is Better
SE +/- 0.24, N = 3
MIN: 27.03 / MAX: 36.83
34.83
MIN: 28.84 / MAX: 39.86
38.17
MIN: 28.41 / MAX: 34.46
33.69
9
18
27
36
45
PyTorch 2.1
Device: CPU - Batch Size: 512 - Model: ResNet-50
batches/sec, More Is Better
SE +/- 0.12, N = 3
MIN: 16.56 / MAX: 20.37
18.38
MIN: 17.86 / MAX: 19.94
19.20
MIN: 15.72 / MAX: 20.02
17.58
5
10
15
20
25
PyTorch 2.1
Device: CPU - Batch Size: 16 - Model: ResNet-152
batches/sec, More Is Better
SE +/- 0.13, N = 3
MIN: 8.32 / MAX: 9.79
8.99
MIN: 8.38 / MAX: 9.61
9.07
MIN: 8.78 / MAX: 10.02
9.42
3
6
9
12
15

pytorch hx

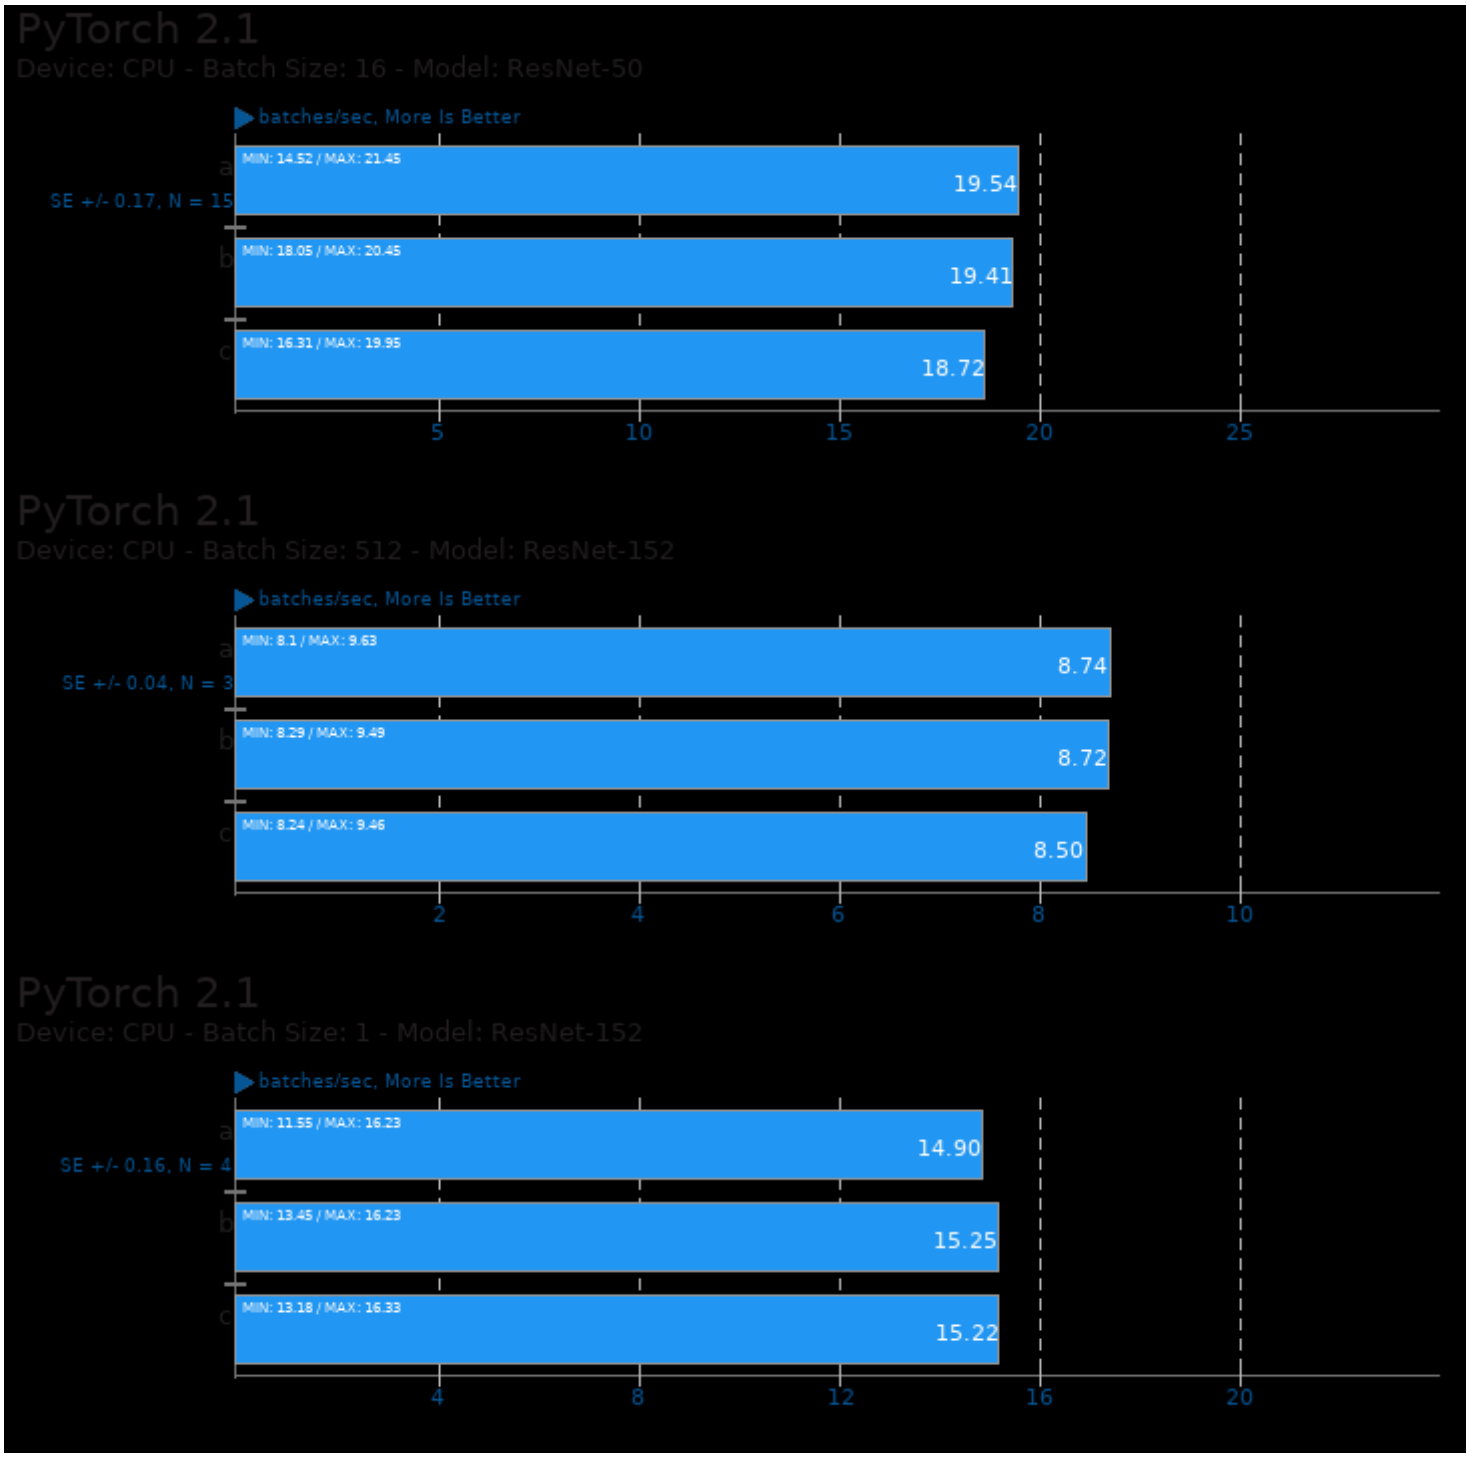
PyTorch 2.1
Device: CPU - Batch Size: 16 - Model: ResNet-50
batches/sec, More Is Better
SE +/- 0.17, N = 15
MIN: 14.52 / MAX: 21.45
19.54
MIN: 18.05 / MAX: 20.45
19.41
MIN: 16.31 / MAX: 19.95
18.72
5
10
15
20
25
PyTorch 2.1
Device: CPU - Batch Size: 512 - Model: ResNet-152
batches/sec, More Is Better
SE +/- 0.04, N = 3
MIN: 8.1 / MAX: 9.63
8.74
MIN: 8.29 / MAX: 9.49
8.72
MIN: 8.24 / MAX: 9.46
8.50
2
4
6
8
10
PyTorch 2.1
Device: CPU - Batch Size: 1 - Model: ResNet-152
batches/sec, More Is Better
SE +/- 0.16, N = 4
MIN: 11.55 / MAX: 16.23
14.90
MIN: 13.45 / MAX: 16.23
15.25
MIN: 13.18 / MAX: 16.33
15.22
4
8
12
16
20

pytorch hx

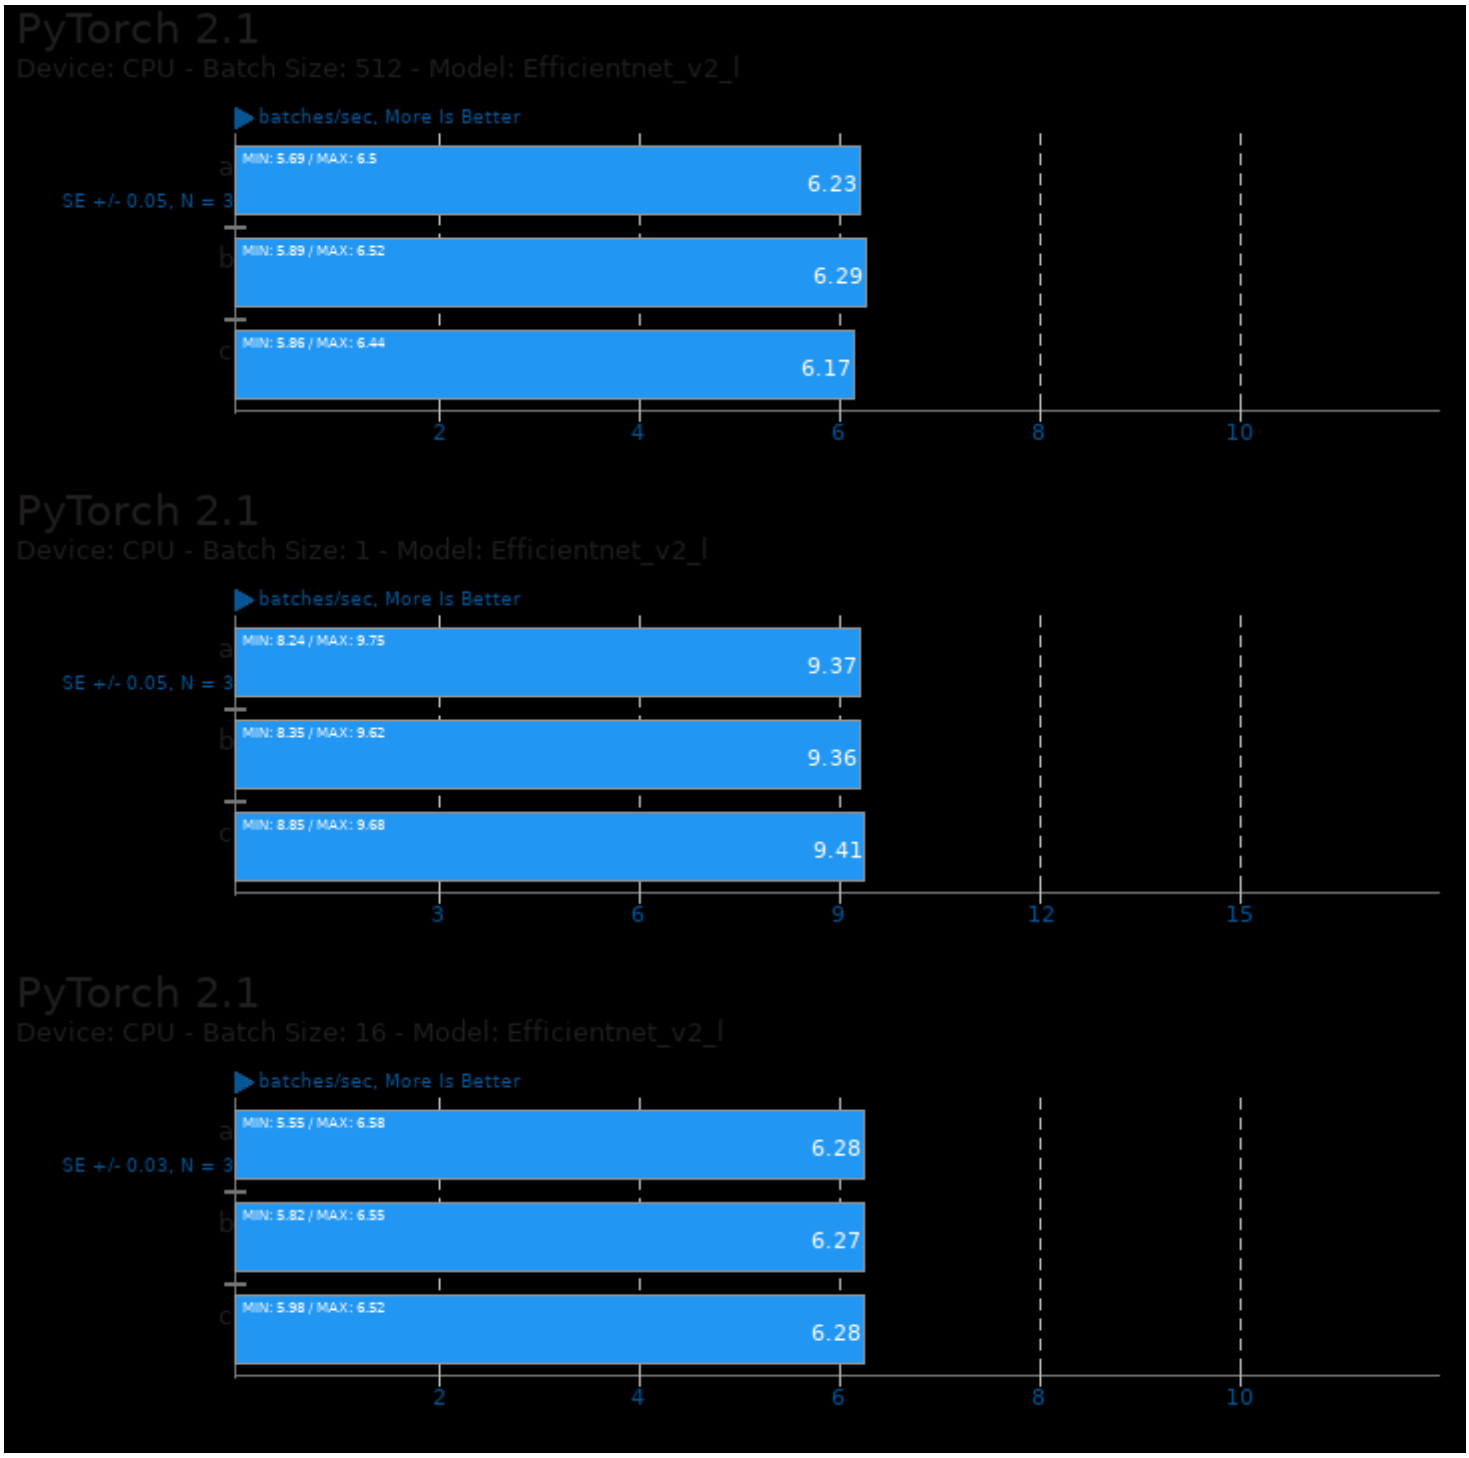

This file was automatically generated via the Phoronix Test Suite benchmarking software on Wednesday, 11 December 2024 20:02.