60 Years of Pushing Performance Boundaries With the Mainframe

Scott Chapman

Enterprise Performance Strategies, Inc.

Scott.chapman@EPStrategies.com

Contact, Copyright, and Trademarks

Questions?

Send email to performance.questions@EPStrategies.com, or visit our website at https://www.epstrategies.com or http://www.pivotor.com.

Copyright Notice:

© Enterprise Performance Strategies, Inc. All rights reserved. No part of this material may be reproduced, distributed, stored in a retrieval system, transmitted, displayed, published or broadcast in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without the prior written permission of Enterprise Performance Strategies. To obtain written permission please contact Enterprise Performance Strategies, Inc. Contact information can be obtained by visiting http://www.epstrategies.com.

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check®, Reductions®, Pivotor®**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM®, z/OS®, zSeries®, WebSphere®, CICS®, DB2®, S390®, WebSphere Application Server®, and many others.

Other trademarks and registered trademarks may exist in this presentation

Abstract (why you're here!)

The modern mainframe traces its lineage back to April 7, 1964 when the IBM introduced the System/360. Businesses were immediately enamored with the radical idea that one shouldn't have to rewrite one's software when upgrading one's computer. However, the other radical part of the announcement was the idea that a single computer line could encompass a 50x range in performance and capacity. The foundation for today's high-performance computing environment was laid in the early days of the System/360 and System/370 and echoes through the decades to enable our modern world.

In this presentation, Scott Chapman will discuss how the IBM mainframe line introduced several performance breakthroughs that are still relevant today. You might even get a few performance tips for your modern mainframe environment along the way.

TL;DR

Today's performance is dependent on yesterday's innovation

AKA: The mainframe has been amazing for an amazingly long time

EPS: We do z/OS performance…

- Pivotor - Reporting and analysis software and services
  - Not just reporting, but analysis-based reporting based on our expertise
- Education and instruction
  - We have taught our z/OS performance workshops all over the world
- Consulting
  - Performance war rooms: concentrated, highly productive group discussions and analysis
- Information
  - We present around the world and participate in online forums https://www.pivotor.com/content.html

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- WLM Performance and Re-evaluating Goals
  - February 19-23, 2024
- Parallel Sysplex and z/OS Performance Tuning
  - August 20-21, 2024
- Essential z/OS Performance Tuning
  - October 7-11, 2024
- Also… please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

© Robert Rogers

Like what you see?

- The z/OS Performance Graphs you see here come from Pivotor

- If you don't see them in your performance reporting tool, or you just want a free cursory performance review of your environment, let us know!
  - We're always happy to process a day's worth of data and show you the results
  - See also: http://pivotor.com/cursoryReview.html

- We also have a free Pivotor offering available as well
  - 1 System, SMF 70-72 only, 7 Day retention
  - That still encompasses over 100 reports!

© Robert Rogers

EPS presentations this week

What	Who	When	Where
60 Years of Pushing Performance Boundaries with the Mainframe	Scott Chapman	Sun 17:00	Neptune D
Introduction to Parallel Sysplex and Data Sharing	Peter Enrico	Mon 13:15	Pomona
Macro to Micro: Understanding z/OS Performance Moment by Moment	Scott Chapman	Mon 15:45	Neptune D
WLM Turns 30! : A Retrospective and Lessons Learned	Peter Enrico	Tue 10:30	Neptune D
PSP: z/OS Performance Spotlight: Some Top Things You May Not Know	Peter Enrico Scott Chapman	Tue 13:00	Pomona
More/Slower vs. Fewer/Faster CPUs: Practical Considerations in 2024	Scott Chapman	Tue 14:15	Neptune D
z16 SMF 113s – Understanding Processor Cache Counters	Peter Enrico	Wed 13:15	Pomona

Agenda

- Obligatory gee-whiz charts
- When were certain technologies introduced and why do we care today
  - Foundational CPU technologies
  - Some significant functionality milestones
- What next? (Scott's Prognostications, Hopes, and Dreams)
- Bonus quotes/pics

If we're lucky, 60 years will take less than 60 minutes!

Personal note

- This was a difficult presentation to build because I kept geeking out on old manuals, photos and specifications
- My initial thought was that this would be much more technical
  - But I don't have all day to talk…
  - Interesting details of obsolete technology may be little more than trivia
    - We probably should learn something "useful"…
  - OTOH… reading some of those early manuals does help give you a sense for why some things are the way they are!
- Instead, I've tried to make this more reflective about how we're still using performance innovations that started decades ago
  - Hopefully you'll find the history interesting
  - Hopefully you'll pick up a few tips pertinent to today's performace

Note about IBM names

- I've probably messed up, misinterpreted, or abbreviated some IBM marketing names
  - Such situations are unintentional (mostly) and I apologize to the IBM marketers who no doubt spent countless hours determining the best name for the many product lines over the decades

Naming things is hard!

Understanding IBM names is harder!

Gee-whiz charts

Time for the clock speed chart

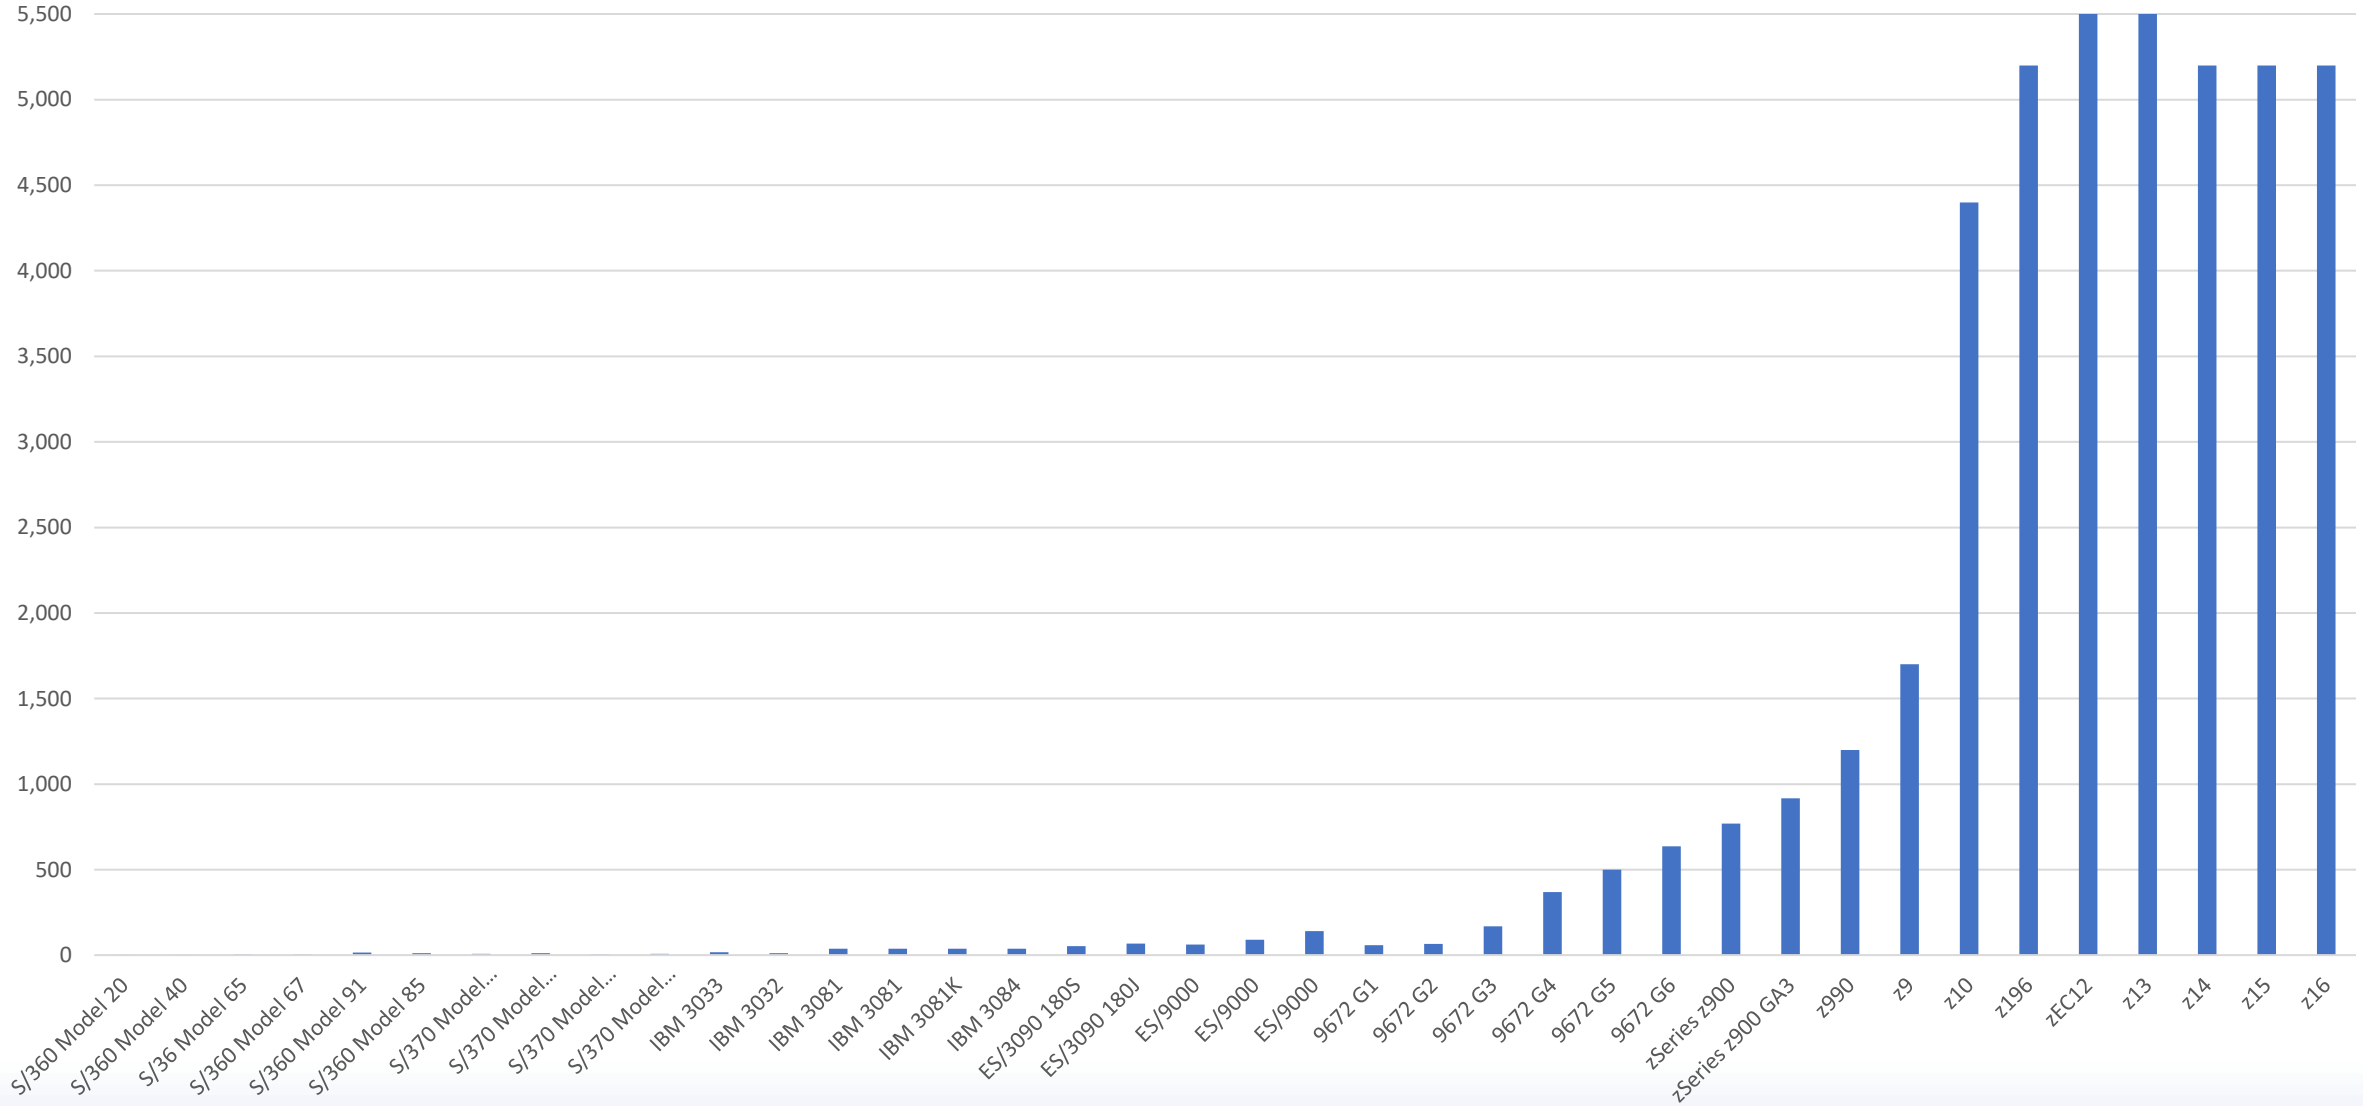

MIPS per CP (uniprocessor config)

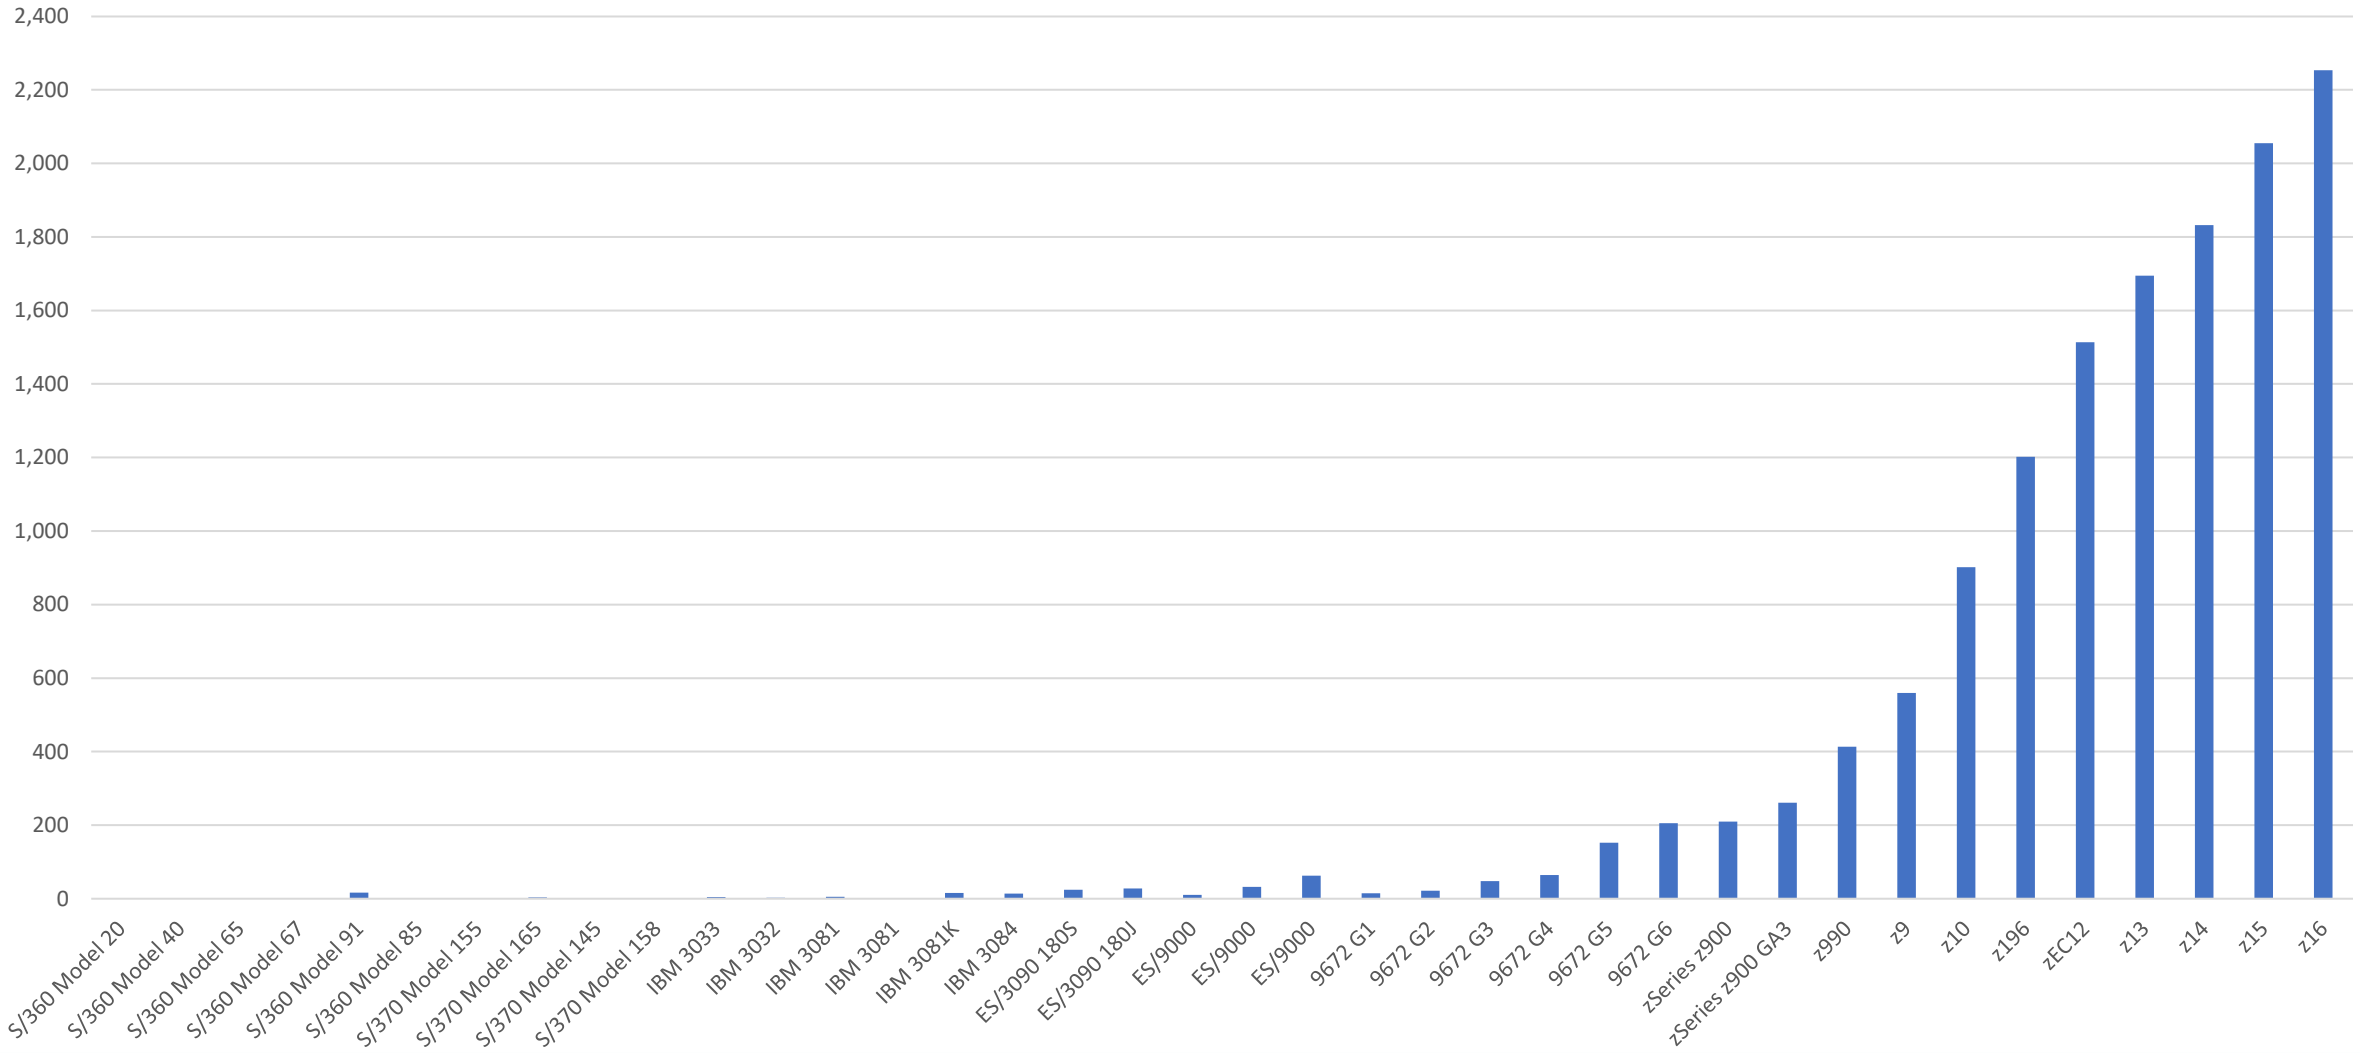

Multiprocessors: Max CPs/machine

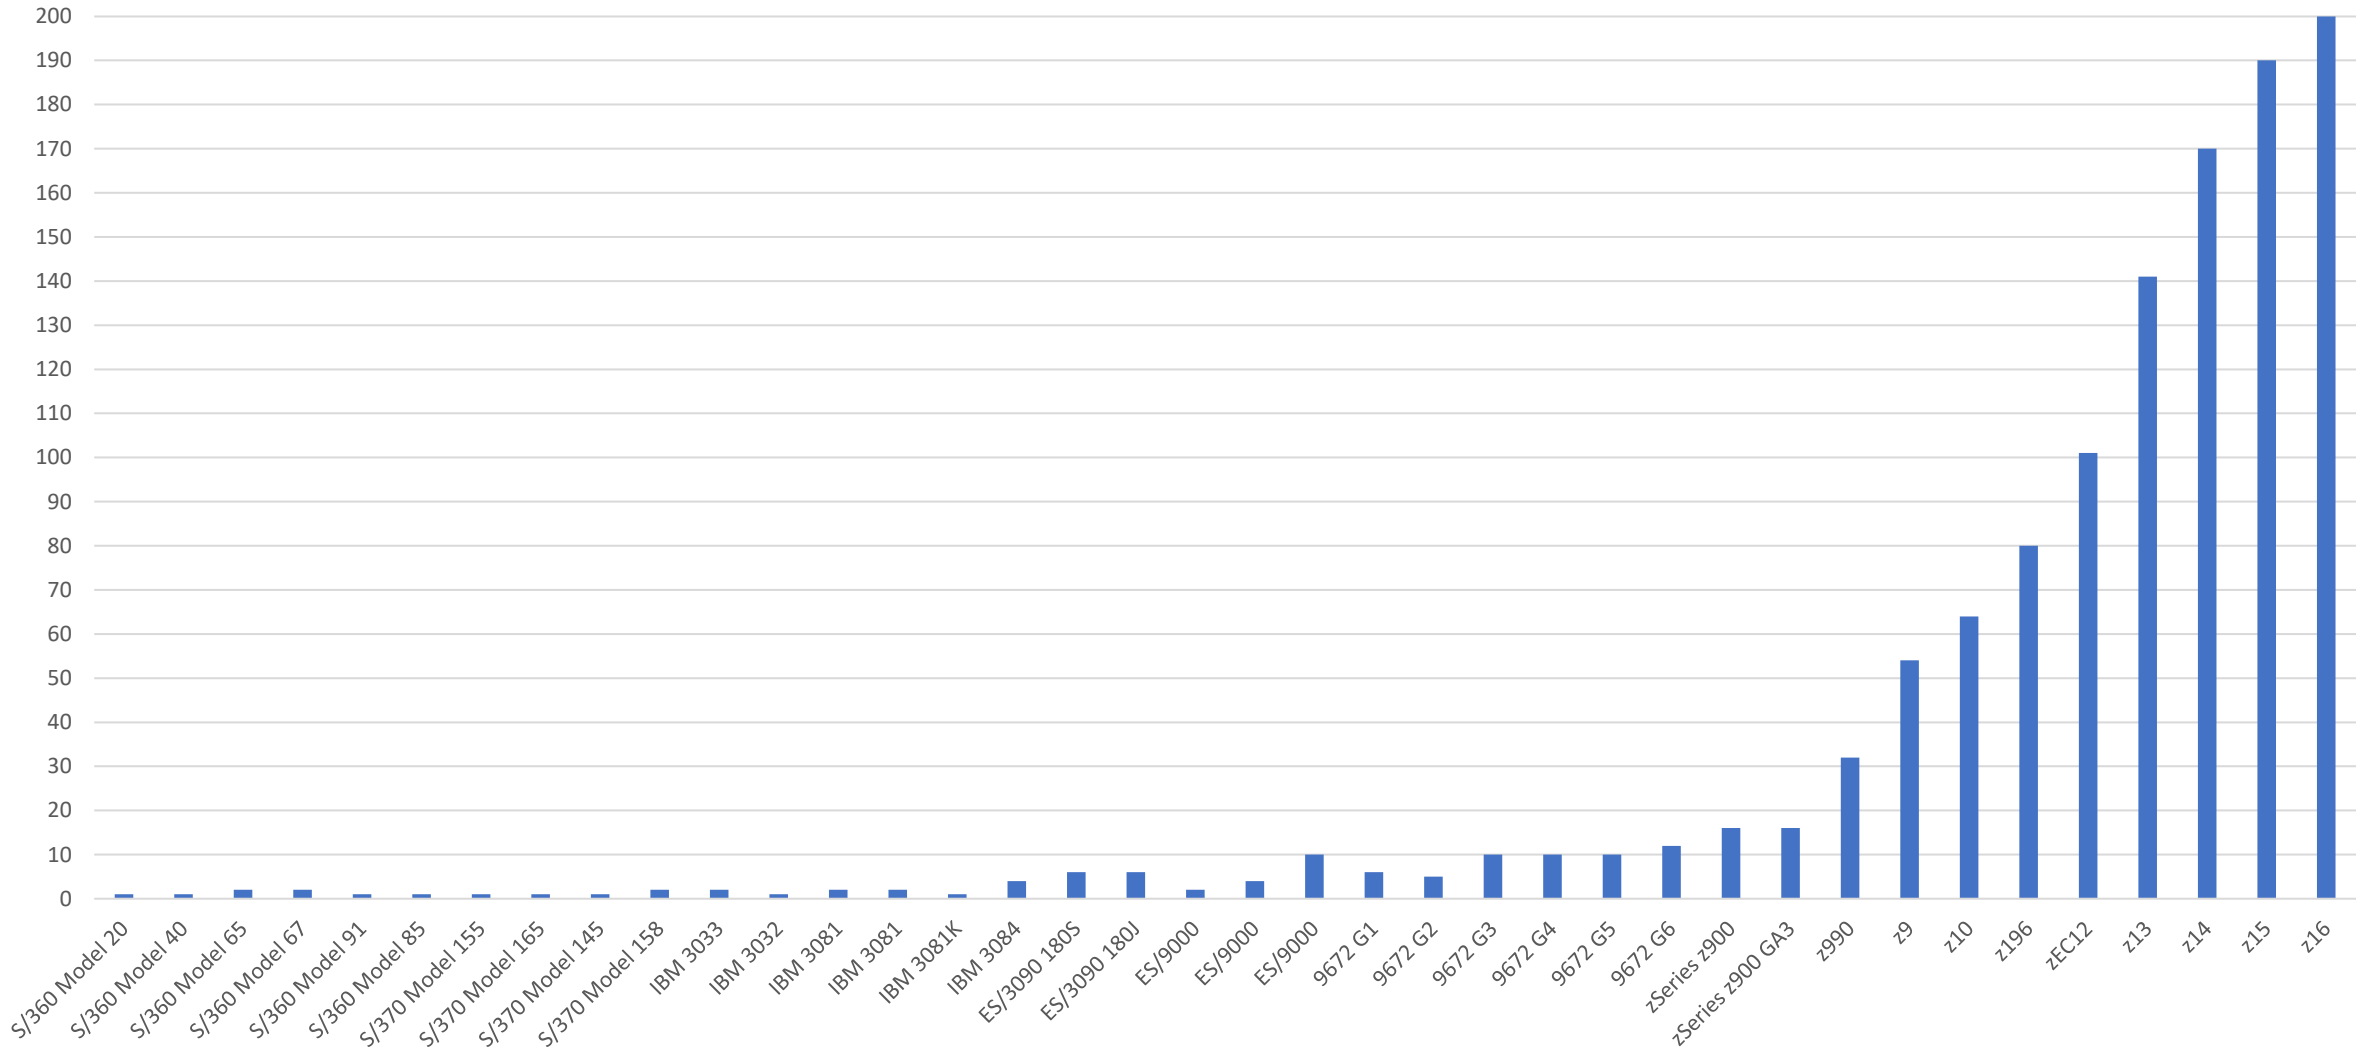

Total Multiprocessor MIPS/machine

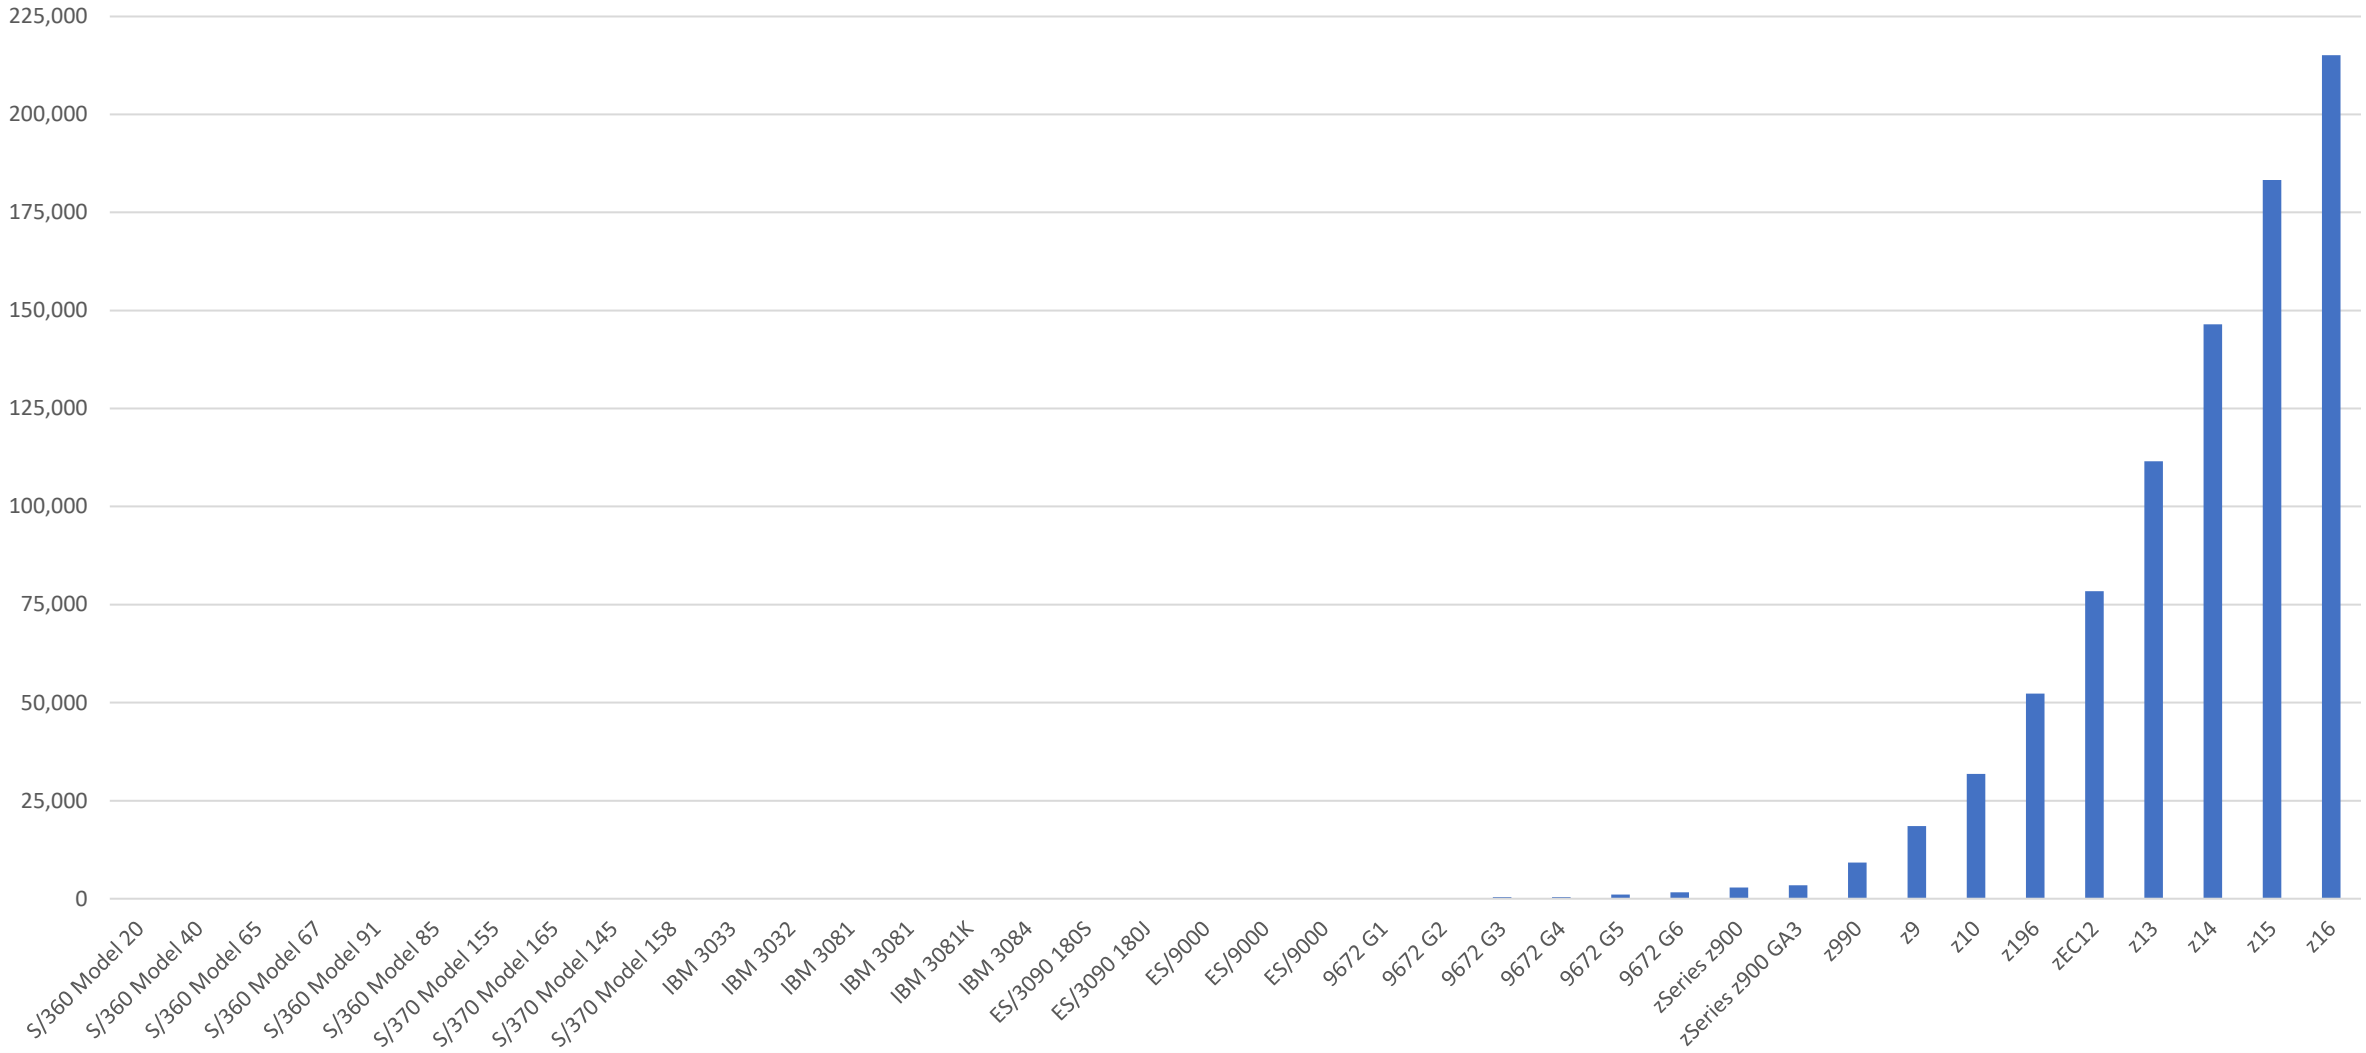

60 Years of IBM Mainframes

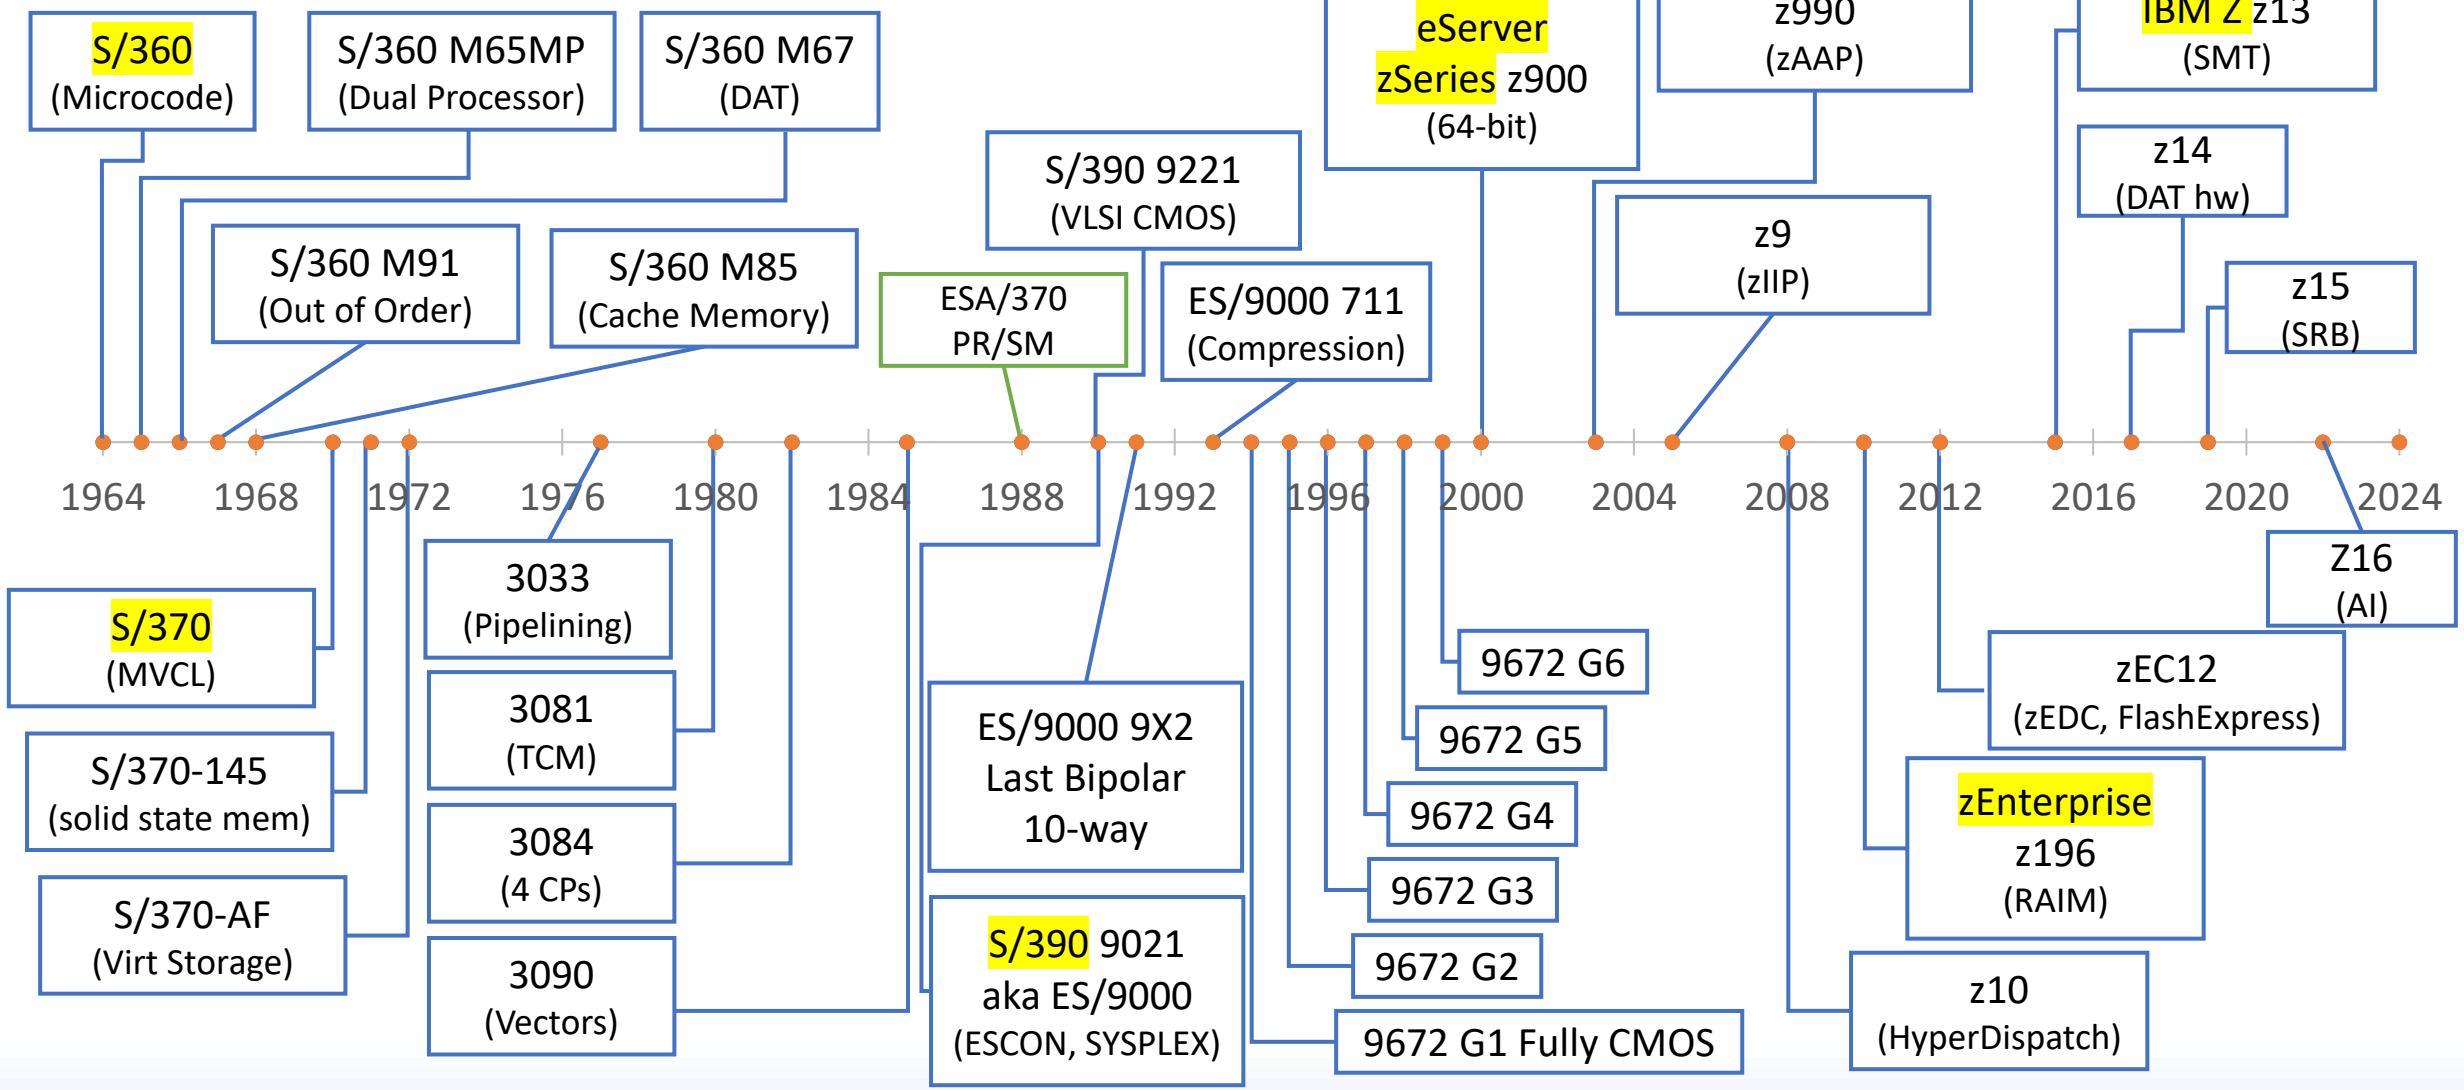

What are all these things?
Why do we care today?

Part 1: Fundamental CPU Technology

1964

S/360
(Microcode)

- Prior to the S/360 compatibility between machines generations (or even models within the same line) was rare
  - New computer? Rewrite all your code!
    - Or at least recompile if you were using that new-fangled COBOL (or FORTRAN)

- The S/360 pioneered the use of microcode to implement the instruction set, making program compatibility ubiquitous across the line
  - Some higher end S/360 models did have hardwired instructions
  - Could emulate older machines too
  - Microcode was in ROS: Read-Only Storage

To shift to 1400 compatibility modes, the customer simply loads in the compatibility microprogram deck.
System /360 Field Engineering Announcement

- Microcode implements one architecture on a (often) simpler machine

Microcode today

- Microcode controls many, many functions of the machine
- Some instructions continue to be implemented in microcode, millicode or even picocode
  - Early CMOS S/390 used microcode to implement S/390 on simpler hardware
  - Faster, cheaper time to market
- Millicode came about c. 9672 G4
  - Runs on same processor as target architecture, like normal programs
  - Implements complex instructions, I/O functions, RAS, virtualization, etc.
  - Millicode has access to hardware facilities not available to normal code
- Picocode
  - Simpler micro-ops needing little decoding, possibly doing >1 op/cycle
  - Dynamic Address Translation apparently done in picocode prior to z14

1965

S/360 M65MP
(Dual Processor)

- The S/360 Model 65 MP was a dual-processor machine
  - Comprised of two 2065 CPUs and 2 to 4 processor storage units (up to 1,048K!)
  - Shared storage and shared I/O
- Could run in Multisystem mode, Mod 65 mode, or Partitioned Mode
  - First two allowed direct CPU to CPU communication
  - Second disabled CPU to CPU communication

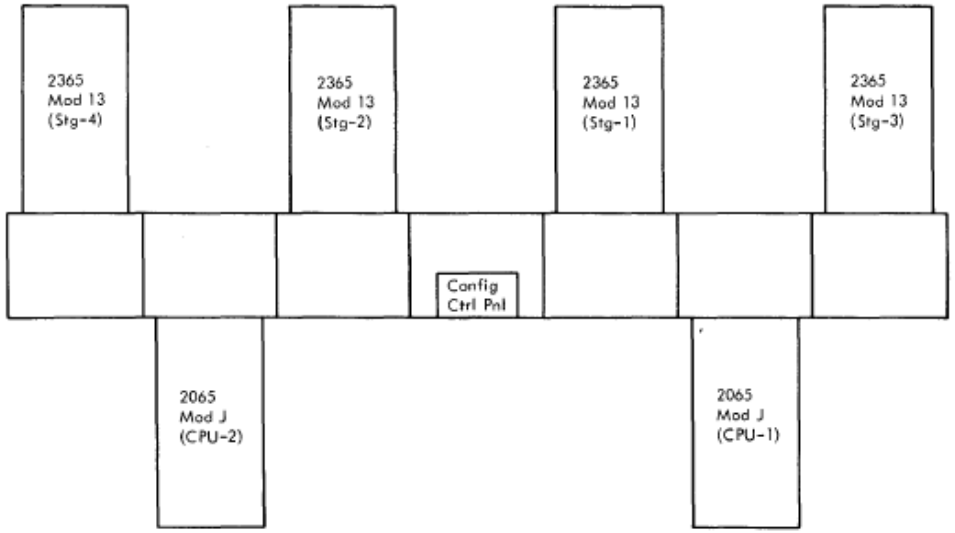

Figure 6. Multiprocessing System Outline Configurations

The extra cable length and priority circuits cause a performance degradation on a shared storage system. This occurs on every storage reference, regardless of the CPU mode of operation. The time added to each storage reference will not exceed the following figures:

	Time per Reference (in ns)	
Storage	CPU1	CPU2
1	30	65
2	65	30
3	30	100
4	100	30

Core storage access time: 8µs (8000ns)

© Enterprise Performance Strategies

Multiprocessing today

- Capacity scaling today is dependent on keeping many CPUs busy
- Most machines have more than 2 CPUs
  - Of course, those CPUs are cores on a chip: (square mms instead of square feet!)
- z/OS really “wants” more than 1 CPU
  - Just because you can run on a uniprocessor doesn’t mean you should
- Still: distance matters
  - You don’t want a CPU accessing memory in a different drawer
  - Keep LPAR sizes small enough (in terms of CPUs and memory) to fit in a single drawer

1966 (shipped)

S/360 M67
(DAT)

- S/360 Model 67 added Dynamic Address Translation
- Separate physical “DAT box”
  - Contained segment and page tables
- Supported virtual memory for the “Time Sharing System monitor” to provide simultaneous multiuser access to multiprocessor system
  - Provided 24 and 32(!) bit address spaces
  - TSS was cancelled twice (’68 and ’71)
  - MTS (Michigan Terminal System) was much longer-lived: I used it in 1987 at RPI
  - CP/CMS was first virtual machine OS, evolving into VM/370 and z/VM

Virtualization Addressing Today

- Everything runs in its own virtual address space
  - Hard to imagine a day when that wasn't the case!
- Dynamic Address Translation still takes CPU time
  - Before the z14 this was quite significant 5-10% of all CPU time was just DAT
    - In some cases even more!
  - More in range of 1-3% on latest machines
- Use large pages to reduce the number of DATs
  - Even though benefit is less than it used to be, it's easy to do
    - Primarily for database buffers and Java JVMs

1967

S/360 M91
(Out of Order)

- The S/360 Model 91 was a scientific machine
  - First one was shipped to Goddard Space Flight Center
  - Went into use by NASA in 1968 after government testing
  - Said to be capable of 16.6 MIPS, at the time the fastest in customer operation
- Its floating point unit was the first to use out-of-order execution
  - Architected by Robert Tomasulo at IBM
  - Was revolutionary for the time
  - Suffered from imprecise exceptions, making it unsuitable for general business use
- Only around 15 Model 91s were built
  - Would not be another out-of-order execution unit in a mainframe until 1991 with the ES/9000 900
  - But then was in-order on CMOS until the z10

Out-of-Order Execution Today

- Common now on all CPU architectures
- One of the advanced techniques that let us continue to get more performance out of CPUs despite relatively stable clock speeds
- Has not been all rosy: Spectre and Meltdown in 2018
  - Side channel attacks based on branch prediction affecting basically all modern CPUs
  - Initial mitigations had significant performance impact
  - Additional similar vulnerabilities found since then

1968

S/360 M85
(Cache Memory)

- Added 16K cache (expandable to 24K or 32K)
- Offsets impact of 1.04µs main storage latency vs 80ns CPU clock
  - That was probably internal latency: likely several ns to get there too
  - 13+ clock cycles to get to main memory is not bad by modern standards
  - Speed of light matters less when your clock rate is that slow!
- Used 64 byte blocks as part of 1K sectors
- Cache hit used 2 clock cycles
- IBM internal testing showed average 96.8% cache hit rate

"…what makes the cache work is the fact that *real programs are not random in their addressing patterns*."
Structural aspects of the System/370 Model 85: II The cache
IBM Systems Journal Vol 7, No. 1 1968

Processor Caches Today

- Processor cache effectiveness is even more key to performance today
  - Speed of light matters a great deal: main memory is much farther away in terms of clock cycles
  - Modern goal is 1 clock cycle for L1 cache access
- Multiple levels of cache, recently:
  - Level 1 – in processor core (~128KB I + 128KB D)
  - Level 2 – dedicated to core (few MB each I + D)
  - Level 3 – shared by all cores on a chip (total 64-256MB)
  - Level 4 – shared by all cores/chips (100s of MBs)
- z16 uses virtual L2-L4 by sharing large (32MB) on-core caches across cores
  - This significant change seems to be working as fairly well

Because of increase in clock speeds and need to keep L1 cache access low, L1 cache size grown less relatively speaking than many

1977

3033
(Pipelining)

- The 3033 added instruction pipelining
  - Earlier models had some pipelining but 3033 extended this capability to many more instructions
  - Said to be 8-stage pipeline
- Improves performance by overlapping work within the processor
  - Still only executing one instruction stream

The 3033 Processor contains an instruction preprocessing function and an execution function that overlap instruction fetching and preparation with instruction execution to improve performance. The instruction preprocessing function is totally controlled by logic circuits and, therefore, can process several instructions concurrently while the execution function is executing a single instruction.

A Guide to the IBM 3033 Processor Complex, Attached Processor Complex, and Multiprocessor Complex of System/370, IBM, 5th Edition 1979

Pipelining, etc. Today

- Processors have only gotten more complicated and more dependent on the performance techniques pioneered decades ago
- Modern high clock frequencies leave little room for error
- Speed of light becomes even more significant limiting factor
- Pipeline flushes cause significant performance impact!
  - E.G. Store into Instruction Stream (updating data within 256 bytes of the executing instructions)
  - Remediate systems (update assembler code) where high SIIS is noted

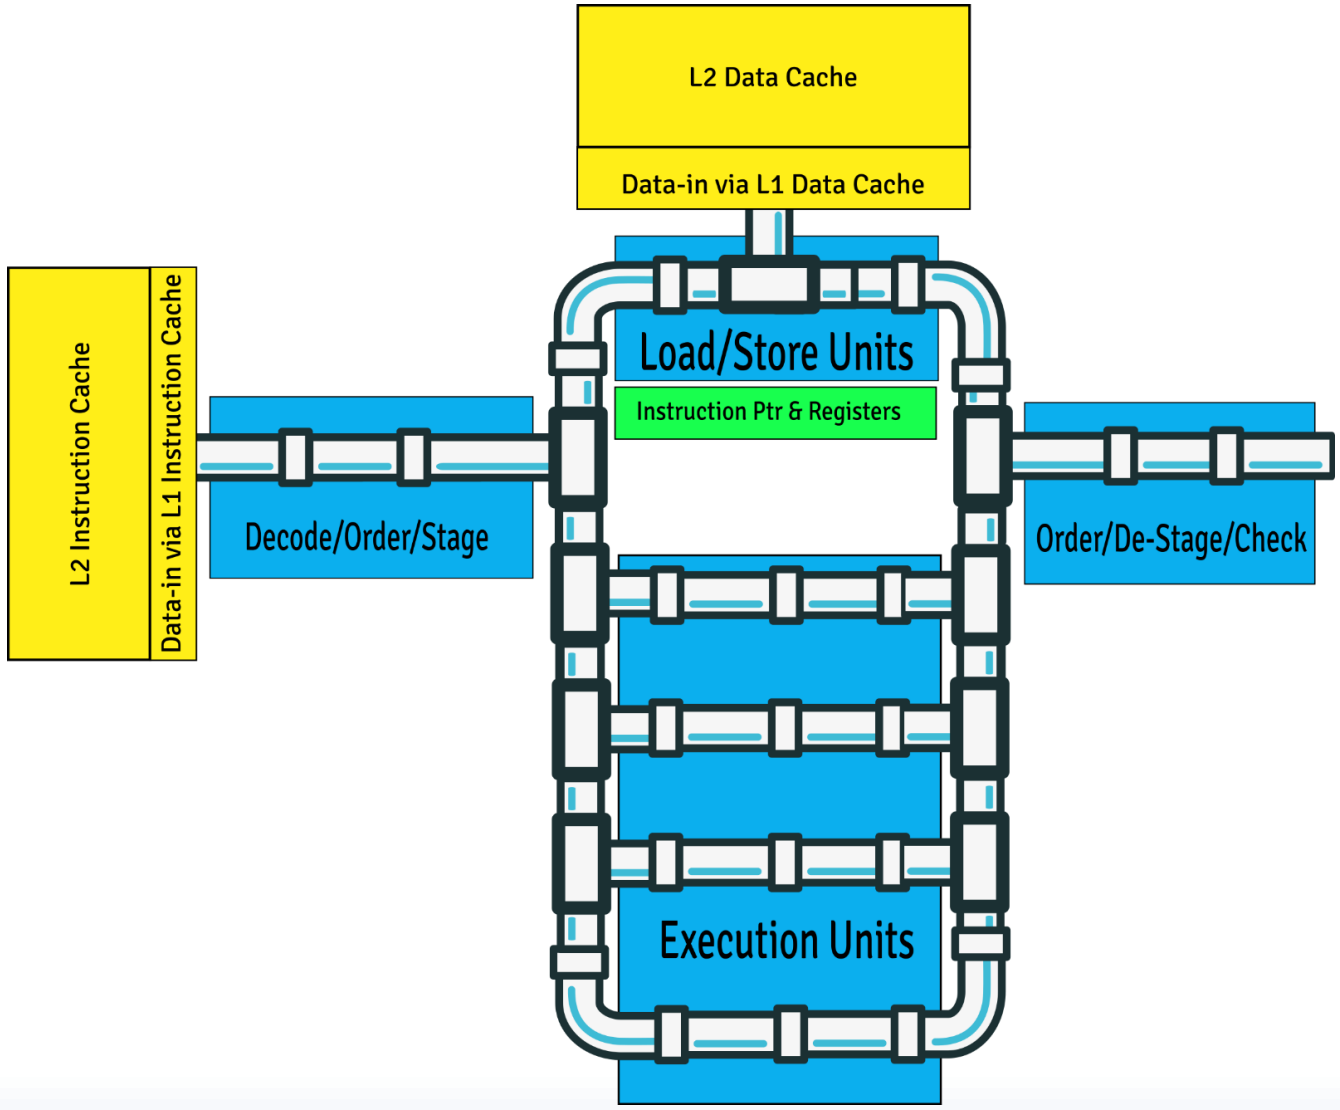

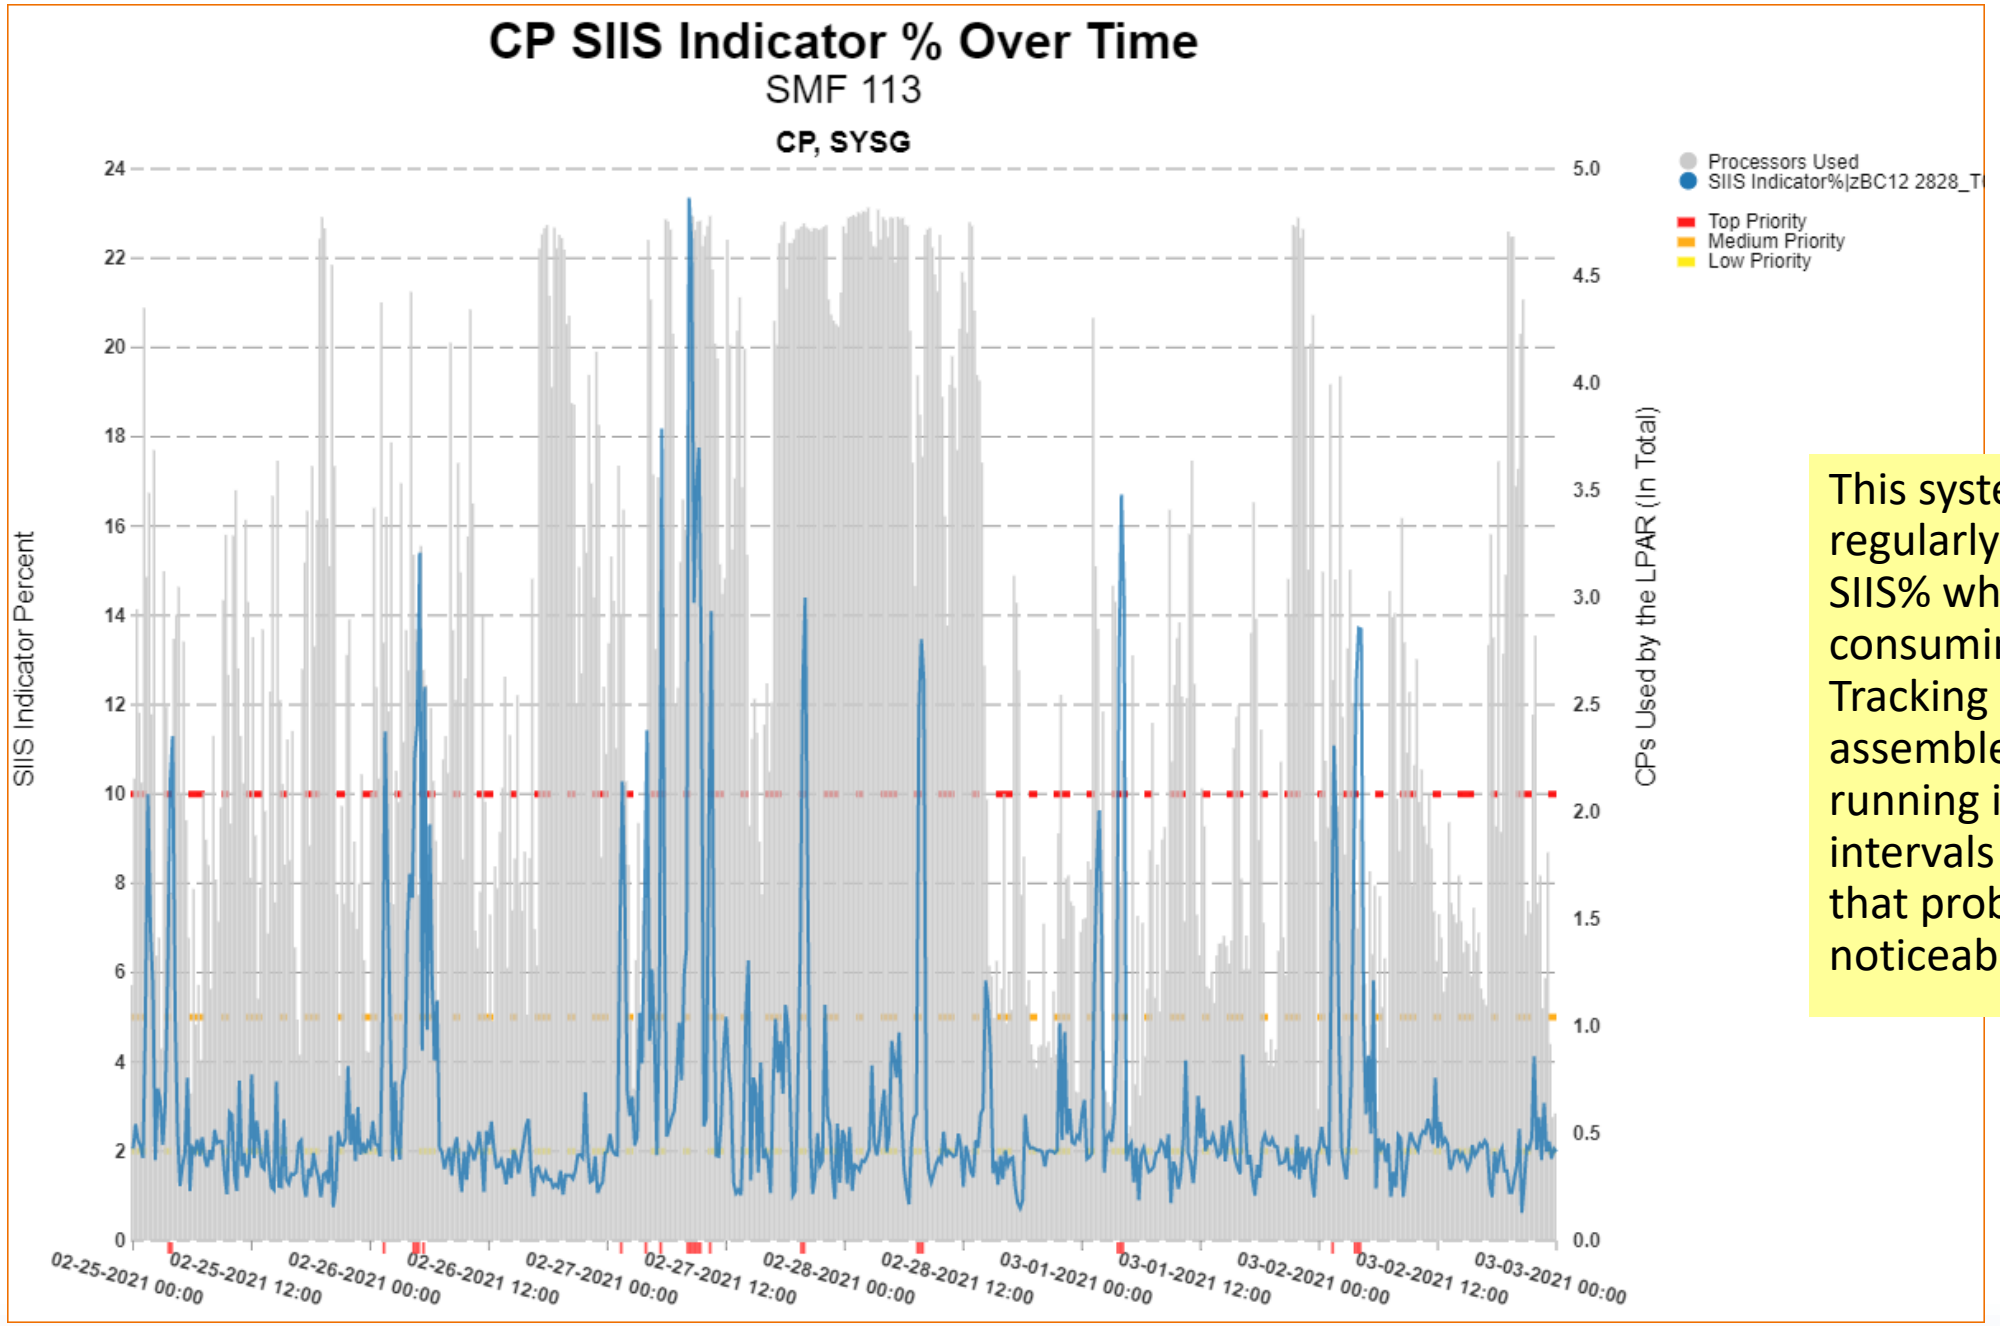

This system seems to regularly exhibit a high SIIS% when the LPAR is consuming multiple CPs. Tracking down what assembler code is running in those intervals and causing that problem may save noticeable CPU.

1980

3081
(TCM)

- Thermal Conduction Model
  - Multi-chip module providing wiring, power, and cooling
- Water cooled via cold plate
  - Up to 118 spring-loaded pistons transfer heat from chips to cold plate
  - Helium filled to reduce thermal resistance around pistons

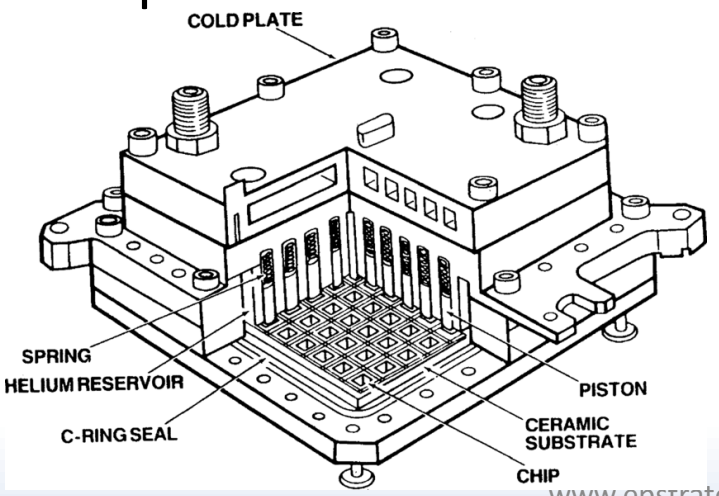

The processor utilizes a new highly integrated circuit technology packaged in Thermal Conduction Modules (TCM).

The TCM is a helium encapsulated module measuring 125 X 134 X 35 mm.

It has 1800 I/O connectors to provide power and to transfer logic signals. Inside are around 100 logic chips.

Heat generated by the circuits is conducted to the outer covering of the TCM, where a cold plate is attached (coolant is circulated through this cold plate).

Those TCM's are mounted on special boards. We have two types of boards, one which can hold 9 TCM's, the other one which can hold 6 TCM's.

On a basic machine we have two 8 TCM boards, two 6 TCM boards and twenty five TCM's (27 full configuration).

Processor Controller is card on board technology.

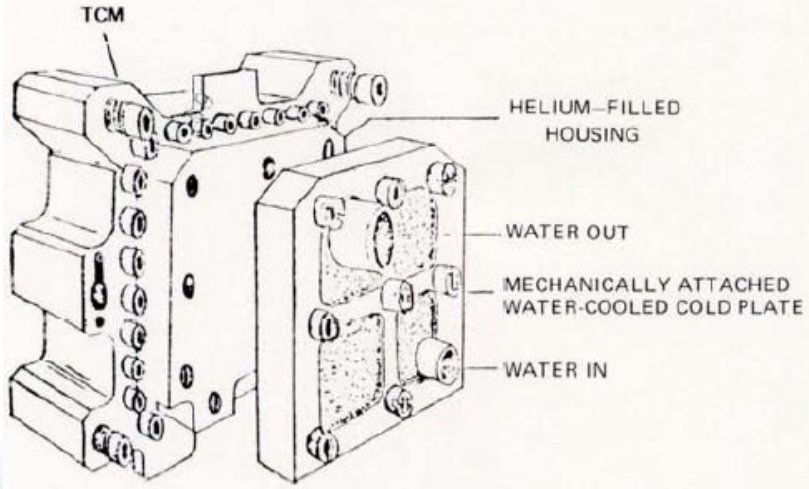

CPU Cooling Today

- Managing processor heat continues to be done with internal closed-loop water cooling
  - “Drain and fill” kit shipped with machine (although now not forced 1 for 1)
- Some PCs even use water cooling
- One of the many problems with increasing clock frequency is heat
  - Usually requires more voltage which leads to more heat more heat
  - Some PC enthusiasts have gotten processors to 8Ghz using liquid nitrogen
    - Probably not practical for “real” use cases
    - Current record as of March 2024 is apparently >9Ghz cooled with… liquid helium!

Yeah, maybe not practical…

https://hwbot.org/benchmark/cpu_frequency/halloffame

Also: they seemed to only run for a second or two.
https://www.youtube.com/watch?v=CzAZxx_aLdo&t=2s

Part 2: Significant Functionality Milestones

1994 Transition to CMOS

- IBM needed to make the transition to a full CMOS architecture
  - Different, cheaper to manufacture, more power efficient transistor type
  - But differential-current switched logic (DCS) using bipolar junction transistors was faster
    - Emitter-coupled logic (ECL) was used prior to S/390 and is largely similar to DCS
- Amdahl and Hitachi were already producing CMOS machines
- New Parallel Sysplex coupling technology was part of the path forward for more, cheaper, slower vs fewer, expensive, faster machines
- New machines were indeed slower but iterated very quickly
  - New generation every year!

Personal note: transitioning from 9672-RY5 to (I think) 9672-R46 was my first exposure to the issues caused by moving to fewer/faster CPs. (Although I didn't fully understand it at the time.)

2000 Transition to 64-bit

- Greatly expanded the address space available to programs to functionally "infinite" (at least for today)
- Was non-trivial to design/implement
- Amdahl and Hitachi didn't make the hardware investment and dropped out of the market
  - Reportedly they combined for 21% of the mainframe market in 1999
  - Both announced they would drop out in 2000

2010 RAIM

- Redundant Array of Independent Memory
  - Basically RAID for memory
- Bit flips do in fact happen, and can ruin your day
- RAIM adds extra layer of protection on top of ECC memory
- Reportedly has been 100% successful in eliminating memory corruption errors
- Note there were functional changes to the z16 memory: now RAIM-like
  - Hopefully this will remain just as robust

2015 SMT

- Simultaneous Multi-Threading only came to the mainframe despite IBM researching it as part of its Advanced Computer Systems project in the 60s
  - Was to design new scientific supercomputer systems
  - Cancelled in 1969
  - Did influence later processor designs
- SMT was relatively late to the mainframe
  - Had been on other platforms for several years
- Due to performance variability and capacity planning difficulties, don't expect to ever see it enabled for any processor where software cost is based on the processor capacity!
- Can be useful solution, but is sometimes over-used
  - Enable SMT on zIIPs when needed, not "just because"

SMT Enablement Decision Flowchart

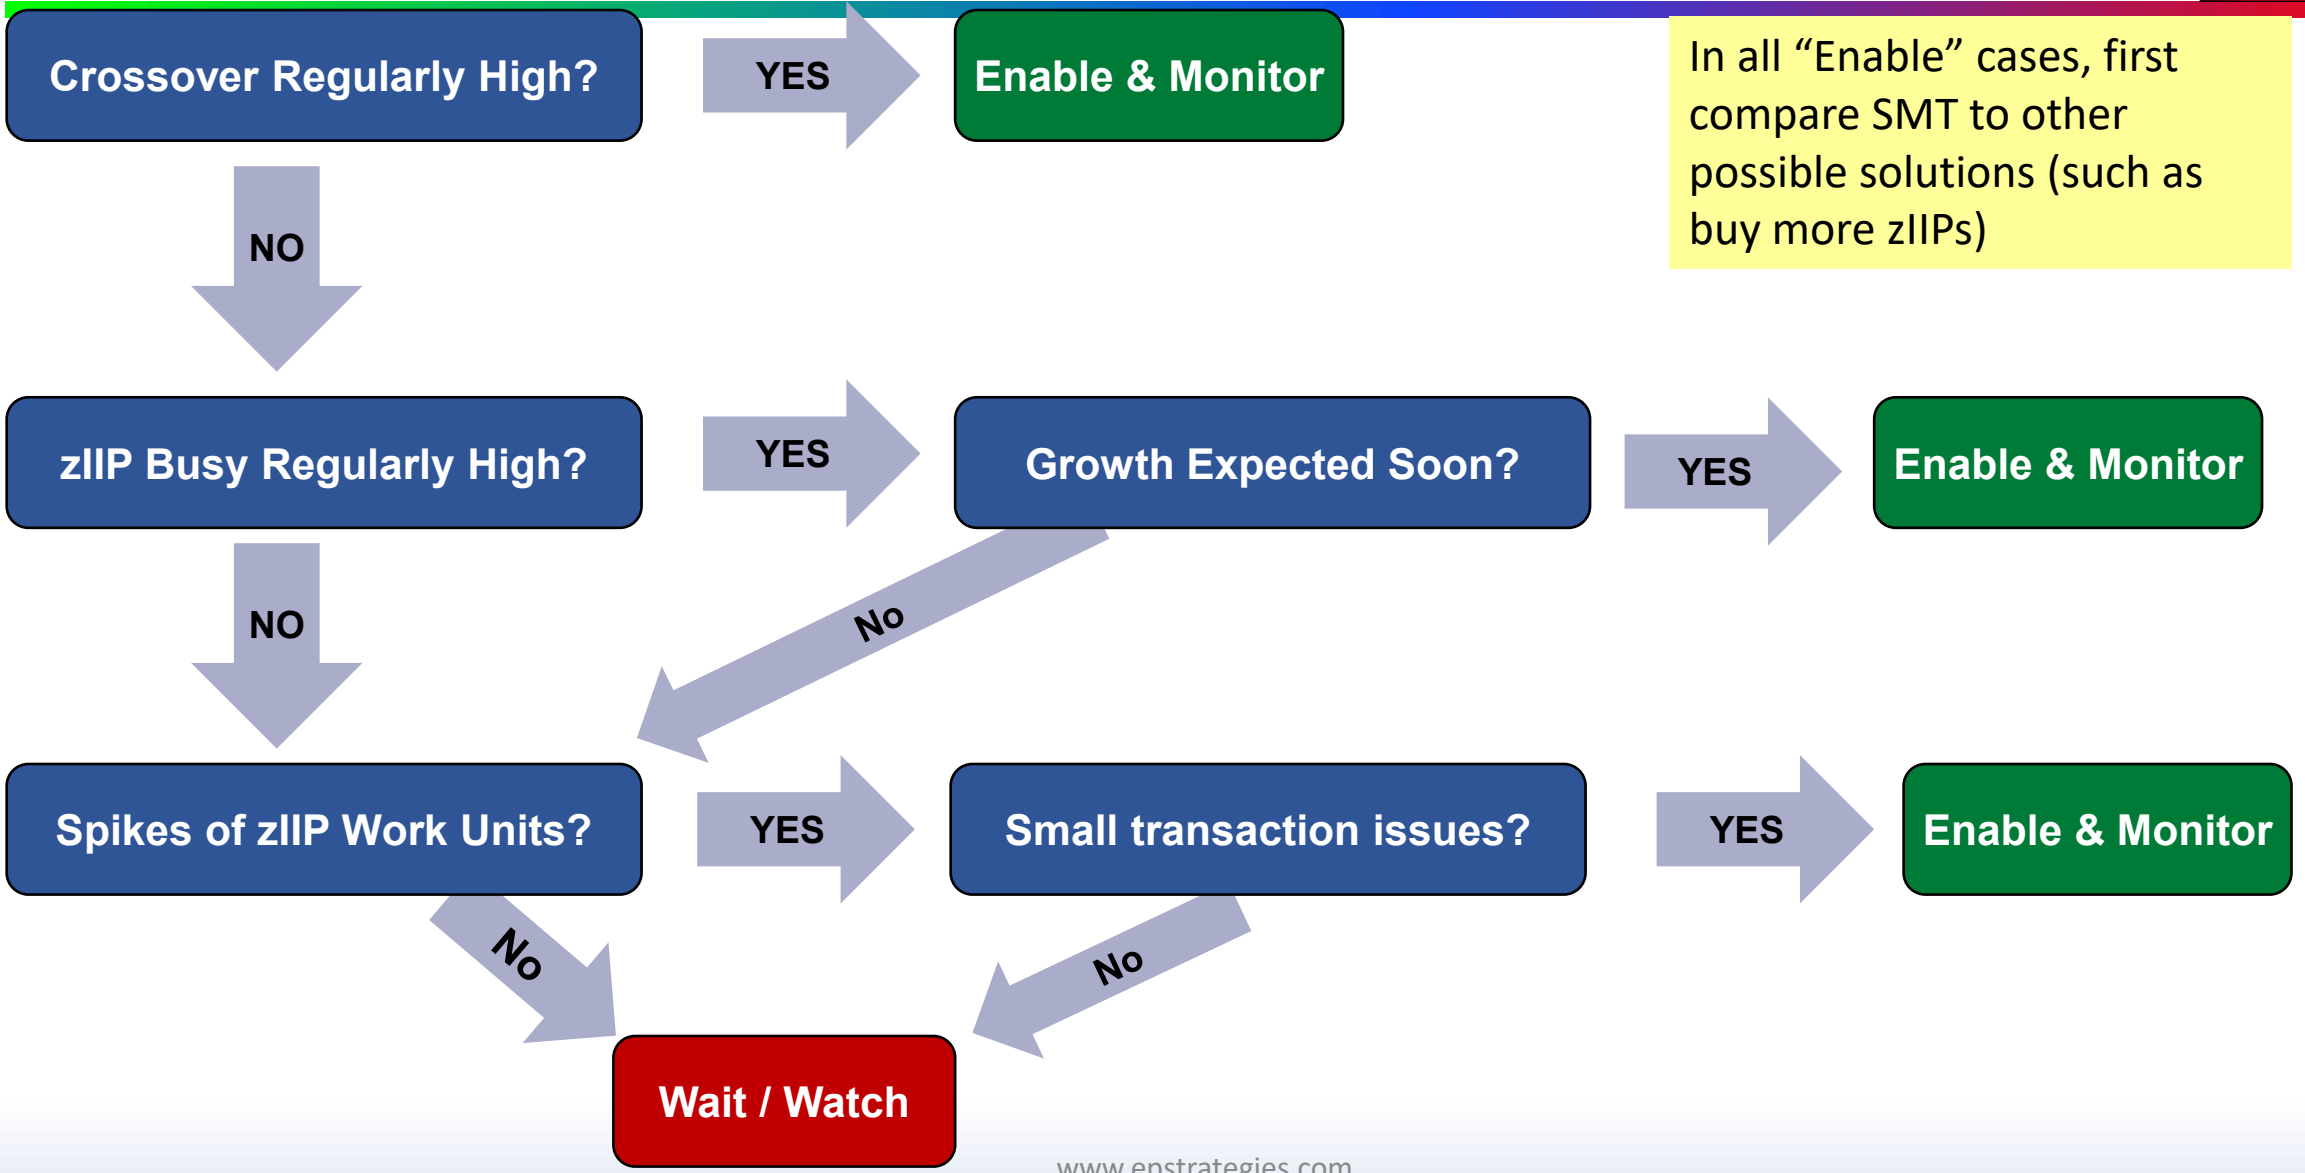

In all "Enable" cases, first compare SMT to other possible solutions (such as buy more zIIPs)

Compression over the years

- In 1993 IBM added compression engine to the CPU, exposed as 2 macros:
  - CSRCESRV – Run Length Encoding, extremely lightweight
    - Used by Db2 and CICS to compress SMF records
  - CSRCMPSC – Ziv-Lempel dictionary based compression
    - Used by Db2, DFSMS, etc.
    - Hardware “accelerated” but still some measurable CPU consumption
- 2012 Added zEDC compression cards
  - Industry standard gzip compression offloaded to hardware card
  - Limited CPU consumption
  - Software cost component involved (minimal)
- 2019 z15 moved zEDC to the chip (in the NXU)
  - No additional hardware costs
  - Still may have to pay for software cost (for most use cases)

Effects of compression

- Compression can often be a net positive for performance by virtue of limiting I/O performed
  - I/O is very slow compared to CPU speeds
  - CPU overhead likely nil to low, and compression ratio can be surprisingly high
- For Db2, pages remain compressed in the BP, allowing same amount of memory to cover more effective data
  - More data in BP may mean less overall I/O which can save some CPU too
- SMF record compression (RLE) in Db2/CICS should remain engaged even if compressing the logstreams via zEDC
  - RLE compression is essentially unmeasurable CPU overhead but can net significant reductions for those records
  - zEDC (dictionary based) will compress more and do so faster when given less data

2022 AI

- I've yet to see any real customer z16 AI performance numbers
  - This does not surprise me—it's still early
- But clearly AI is here for the foreseeable future
  - How customers will use it on the mainframe and to what degree, remains to be seen
  - There's definitely interest out there though
- I'm a fan of the mainframe having the HW and SW to leverage AI models
- I'm a bit wary of z/OS as an AI-infused operating system to simplify operations
  - AI can clearly be a "force multiplier" to improve productivity across a variety of roles
  - I'm worried about using just system measurement data to drive decisions
    - Significant business and institutional knowledge exists outside the realm of the system
    - Trying to run everything as fast as possible is not necessarily what the business needs

What next?

Scott’s Prognostications, Hopes, and Dreams

Easy Expectations: More, Faster

- Per-CPU capacity will continue to grow for at least a few generations
  - Gains will likely be more incremental though
- Overall max system size will continue to increase
  - Max CPUs per drawer and CEC will likely continue upward
- Max memory will continue to grow (but maybe not with every new generation)
- Most interesting changes will be what new features get added to the chips
  - Maybe more AI processors?

Tactical hopes

These address IBM marketing decisions, not hard technical issues!

- More than 4 “speed” tiers is almost becoming a necessity in the high-end
  - 7xx engines are too fast for many customers
  - Steps between tiers too large in some cases too
- Would be nice to see an increase in the number of CPs available on the low-end machines from 6 to maybe 10 (or more!)
  - Would make those machines an option for more sites
  - Would potentially solve the first issue for some sites
- Increase minimum memory on the low-end machines from 64GB to at least 128GB, preferably 256GB
  - Of course, at no cost increase ☺ (should be minimal cost to IBM)
  - Positions those machines for one of my dreams…

Dreams (Technical Innovation)

- More tightly integrated storage, in the frame (would not be the first time!)
  - PCIE connection to internal NVME storage?
- z/OS becomes the in-memory operating system
  - Basically, leverage large memory to eliminate 99% of I/O
  - Can do much of this today in Db2, but need OS support to do lower-level caching
    - E.G. Cache in the channel subsystem? (Multiple system coordination would be “interesting”)
    - Stepping stone: make all I/O zHyperLink (or PCIE?) and increase DS8K cache beyond 4TB?
- Live LPAR move
  - Like z/VM can do, but with PR/SM
  - I think GDPS can do the I/O part today, but the LPARs are re-IPLed
  - Automation to move based on environmental changes (e.g. loss of utility power)
- Make memory non-volatile
  - Could possibly do today with combination of battery-backup and internal flash memory for persistence
  - New persistent memory technologies also are being researched
  - Still larger L4 (5?) cache could offset some of the impact of slower persistent memory

Dreams (Non-technical)

- Multiple GP processor speeds per CEC
  - Technically is happening today with System Recovery Boost
  - Could help significantly where there are “small” and “large” LPARs on same CEC
    - I.E. run PROD, DEV, TEST at different processor speeds, but each with multiple LPs
    - As opposed to today, where small LPARs often have tiny share of a couple of LPs
  - Capacity measurement for software pricing more complicated, but doable
- Combine the high- and low-end machines into one with 26 capacity options
  - Seemingly technically easy, but probably eliminates CPU binning benefits
- Blow up the software pricing model and do away with sub-cap engines
  - More, faster for everybody!
  - Yeah, right! We can dream though?

Bonus quotes and pictures

Quotes

1981: Before the 3081, there was a veil of uncertainty over the mainframe market; now, demand for top-of-the-line systems is revitalized.
Datamation July 1981: The CPU Market: A Survey

1982: There were a couple of TCMs that had a wireability problem, but these TCMs were not particularly dense from a circuit count point of view. The problem was that data and address busses to arrays blocked channels that other wires needed to use. This was caused by an arbitrary choice made by the wiring program in the initial x-y direction of the routed wire. This was corrected by changing the sequences in which the networks were fed to the wiring program, which forced the wire router to choose channels that were different from those initially used.
IBM 3081 System Overview and Technology, Clive A Collins, IBM

1990: All the 9021s will use new power units – for the first time based on 230-volt 60-cycle three-phase current rather than the old 415-volt (sic: I think they mean 400Hz) technology.
https://techmonitor.ai/technology/the_forecast_shape_of_the_ibm_370_mainframe_line_for_the_1990s

1997: Ending four years of embarrassment, IBM has at last been able to produce a CMOS microprocessor for its S/390 mainframes that out-performs, on a processor vs processor basis, IBM's traditional bipolar technology.
Specifically, the high-tuned version of the new chip (appropriately named the Turbo) is said to howl away at 63 MIPS, compared with about 60 MIPS for the most fleet-footed mainframe dinosaurs. This contrasts with 50 MIPs for the G3 series, introduced just nine months ago.
https://www.hpcwire.com/1997/06/13/ibms-new-s390-cmos-chip-beats-bipolar-at-last/

2000: "IBM's [64-bit] architecture is very proprietary and requires a significant investment [to emulate]" Stone said. With mainframe sales projected to dwindle in the next few years and demand for 64-bit mainframes likely to remain low for some time, the investment wasn't worth it, she said.
Carol Stone, VP, Amdahl, Computerworld:
https://www.computerworld.com/article/2588995/amdahl-gives-up-on--mainframe-business.html

S/360 Planning

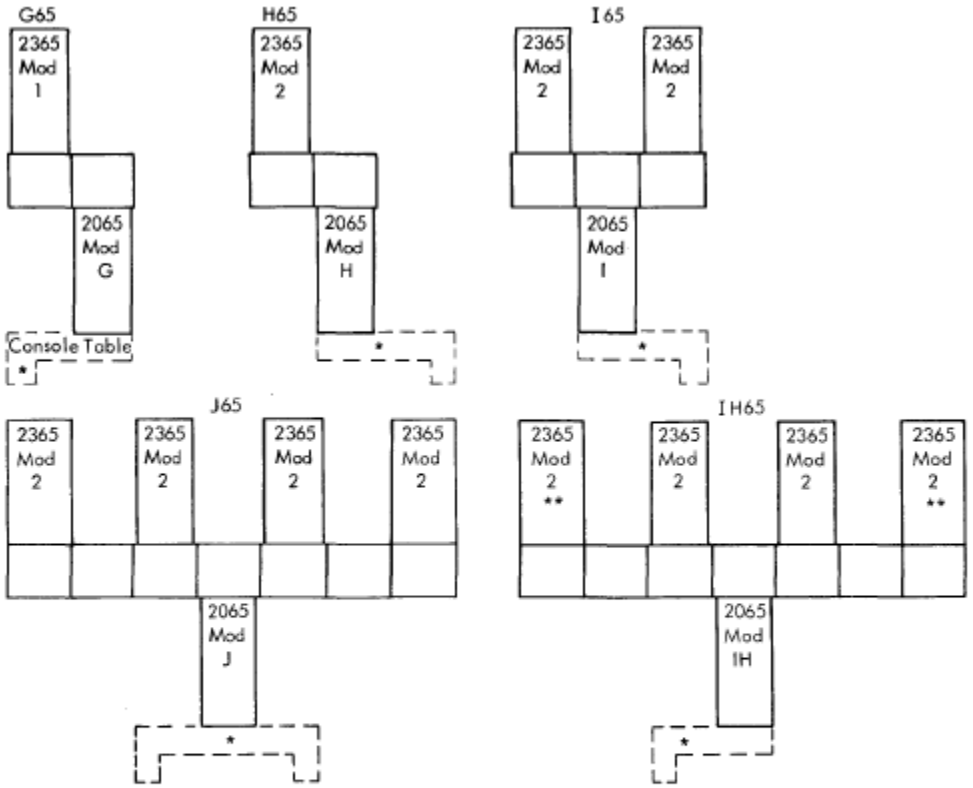

*L-shaped console table may bend to either the left or the right (option). U-shaped console table, as shown for Model J65, is provided on any model that attaches two 1052-7's.

**Third 2365 Mod 2 can be in either end position.

Figure 2. 2065 Processing Unit and 2365 Processor Storage Configurations

S/360: console table included!

MOTOR GENERATOR (REMOTE)
FOR SYSTEM/360 MODEL 85 (60-HZ INPUT)

PLAN VIEW

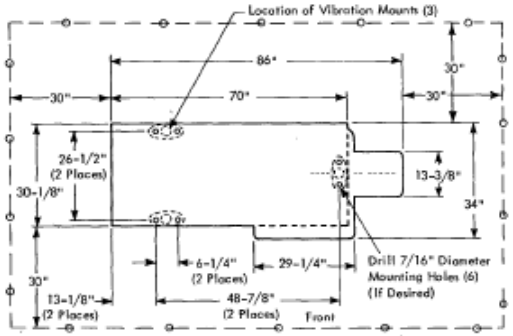

Notes:

The installation and maintenance of the motor-generator (including starter) unit will be the responsibility of the customer.

At time of installation:

1. An overvoltage circuit is provided in the motor-generator regulator. This must be adjusted to remove generator output when the 415-Hz line voltage reaches 220 ± 2V (rms).
2. The generator output voltage must be set so that the voltage measured by the meter located on the power distribution unit (2085 frame 14) reads between the center and upper scribe marks.
3. Consult motor-generator manufacturer's instruction manual for further installation procedures and maintenance.

SPECIFICATIONS

Dimensions:

	F	S	H
Inches	86	34	53
(cm)	(218)	(86)	(135)

Service Clearances:

	F	R	Rt	L
Inches	30	30	30	30
(cm)	(76)	(76)	(76)	(76)

Weight: 4,200 lb (1.950 kg)

Heat Output (Max):

208/230V: 86,000 BTU/hr (21.700 kcal/hr)
440V: 102,000 BTU/hr (25.750 kcal/hr)

Power Requirements:

Input:
Induction Motor–200 hp, 208/230V or 440V, 60 Hz, 240A full load, NEMA design B, code F, 40°C maximum ambient, dripproof enclosure

Output:
Synchronous Generator–175kVA, 208V, 3 phase, 415 Hz, 485A full load, 70°C temperature rise, dripproof enclosure

That's a lot of power!

Possible power supply answer

Why 400Hz power? Building reliable power supplies for such high power was not easy at the time, and CDC took an elegantly brutish approach. First, using 3-phase power and rectifying it can greatly reduce the filtering requirements because the ripple is at 6 times line frequency and thus much easier to filter. Also, the raw ripple doesn't even go down to zero volts and back like single phase rectified AC. Second, one can reduce the filtering requirements as well as transformer sizes by using 400-Hz power instead of 60-Hz. Such power is often used in aircraft and some military applications, so components were available. Including the three 400-cycle motor-generators we had to install in the basement (one per mainframe and one spare). This now meant you could build 100-ampere power supplies with basically three components (transformer, bridge rectifier, filter capacitor). Crude compared to switching power supplies that came along later, but there was no question of reliability with 3 simple parts.
https://web.archive.org/web/20060906162720/http://w3.uwyo.edu/~jimkirk/cyber_era.html

DESIGN CONSIDERATIONS FOR POWERING 415 HZ COMPUTER SYSTEMS

Thomas M. Gruzs
Member IEEE
Liebert Corporation
P. O. Box 29186
Columbus, OH 43229

Since the majority of commercially available electrical power in the world today is either 50 or 60 Hz, why do some mainframe computer systems require 415 Hz input power? 415 Hz power must offer some advantages to the computer vendor over 50 or 60 Hz power. First, 415 Hz power supply components are significantly smaller—the transformers and inductors are smaller for the same power capacity and rectified 415 Hz requires significantly less filter capacitance. Additionally, since 415 Hz power is not commercially available, some form of frequency conversion equipment, such as a motor–generator set or static frequency converter, must be used and is often relied upon to provide power conditioning or voltage pre–regulation. The 415 Hz computer power supplies can be designed with a narrower input voltage specification than 50 or 60 Hz power supplies which are applied to commercial power systems (without power conditioning) where wider voltage variations and other power disturbances are to be expected.

Instructions timings from S/370

Instruction	*Format*	*Op Code*	*Mnemonic*	*Time in Microseconds*
add	RR	1A	AR	1.373
add*	RX	5A	A	2.385
add decimal	SS	FA	AP	$8.757 + 0.744\ N^1 + 1.375\ NWBL^1 + 0.540\ NWBL^2 + T^1$ $(3.044 + 0.451\ N^1 + 1.127\ NWBL^1)$
add halfword*	RX	4A	AH	2.949
add logical	RR	1E	ALR	1.373
add logical*	RX	5E	AL	2.138
add normalized (extended)	RR	36	AXR	12.134#
add normalized (long)	RR	2A	ADR	7.162
add normalized* (long)	RX	6A	AD	8.265
add normalized (short)	RR	3A	AER	5.663#
add normalized* (short)	RX	7A	AE	6.737
add unnormalized (long)	RR	2E	AWR	6.524
add unnormalized* (long)	RX	6E	AW	7.627
add unnormalized (short)	RR	3E	AUR	5.513
add unnormalized* (short)	RX	7E	AU	5.342#
and	RR	14	NR	1.935
and*	RX	54	N	2.700
and	SI	94	NI	2.397
and	SS	D4	NC	$N \leq 4 = 6.437 + 1.148 NWBL^1 + 0.540\ NWBL^2 + 0.405N$ $N > 4 = 6.740 + 1.148 NWBL^1 + 0.540\ NWBL^2 + 0.203N$
branch and link	RR	05	BALR	$1.682 + 0.874F^1$
branch and link*	RX	45	BAL	2.399
branch on condition	RR	07	BCR	$0.872 + 0.875F^1$
branch on condition*	RX	47	BC	$0.917 + 0.875F^1$
branch on count	RR	06	BCTR	$1.074 + 1.078F^1$
branch on count	RX	46	BCT	$1.369 + 0.873F^1$
branch on index high	RS	86	BXH	$2.469 + 0.875F^1$
branch on index low or equal	RS	87	BXLE	$2.469 + 0.875F^1$
compare	RR	19	CR	1.578
compare*	RX	59	C	2.441
compare decimal	SS	F9	CP	$8.577 + 0.451N^1 + 0.789NWBL^1$
compare halfword*	RX	49	CH	2.949
compare logical	RR	15	CLR	1.373
compare logical*	RX	55	CL	2.138
compare logical	SI	95	CLI	1.992
compare logical	SS	D5	CLC	$N \leq 4 = 3.494 + 0.540\ (NWBL^1 + NWBL^2) + 0.405N$ $N > 4 = 3.994 + 0.540\ (NWBL^1 + NWBL^2) + 0.203N$

Small portion of S/370 Model 145 instruction timings

Performance reporting from 1977

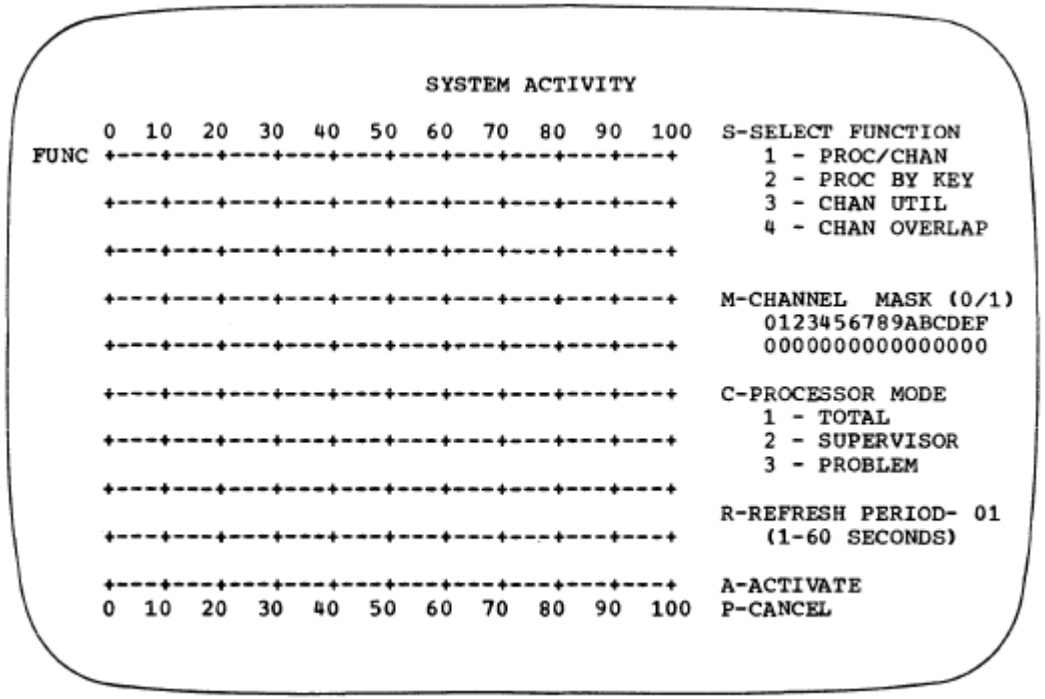

Figure 20.25.5. The system activity frame

IBM 3033 System Activity

Inside an ES/9000 TCM

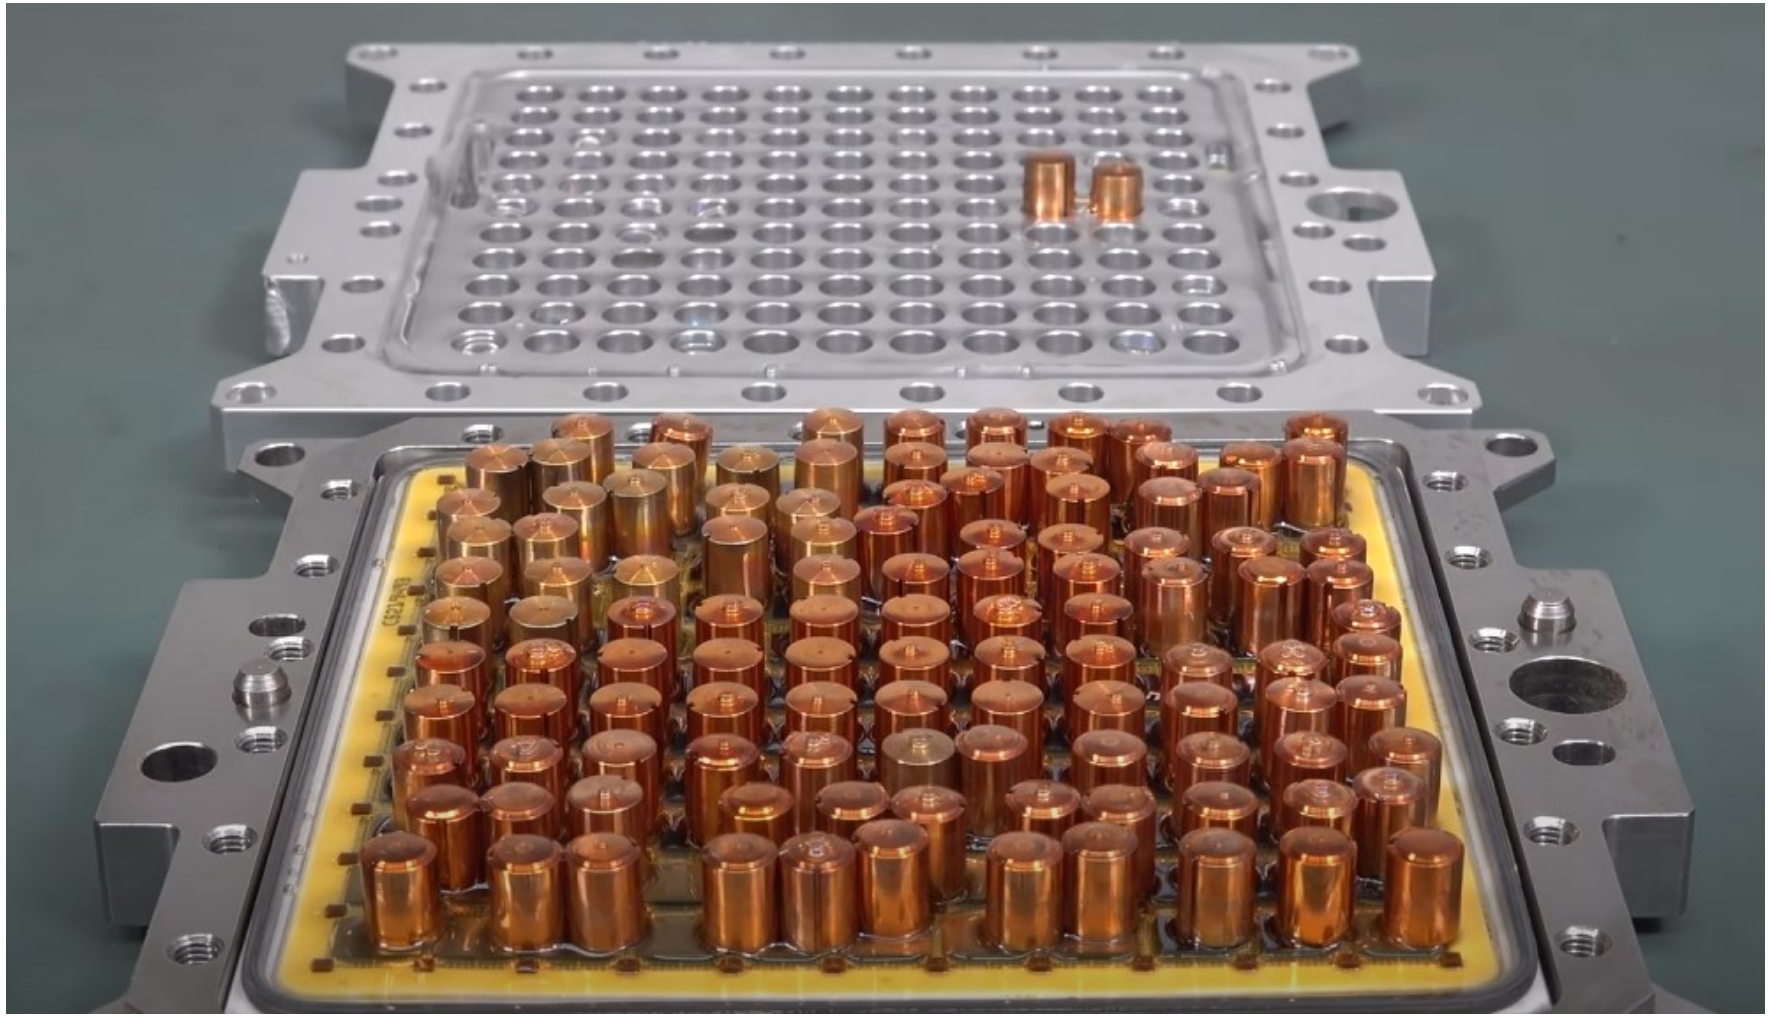

https://www.youtube.com/watch?v=xQ3oJlt4GrI

Questions?

AKA: Why didn’t you talk about…