z/TPF V1.1

2013 TPF Users Group

Title: Performance Topics

Robert Blackburn Ph.D.

AIM Enterprise Platform Software
IBM z/Transaction Processing Facility Enterprise Edition 1.1

HW on floor can have variable MIPS

- **Without an IPL TPF will adjust to the increased/decreased MIPS**
- **As an example with 15 CPs(current maximum EC12) we get 4 different MIPS ratings with no IPL**

		scaling factor ((n+1)/n)
2817-415	77.9 ITRR	----------------
2817-515	186.7 ITRR	2.38
2817-615	238.5 ITRR	1.28
2817-715	365.9 ITRR	1.53

- **Talk to IBM account team**
  - Conditions and HW/SW charges will apply
  - Designed to handle seasonal peaks
  - Infrequent adjustment

2G to 4G for C programs

- **Typically C programs (in 64-bit CRPAs) reside at addresses > 4G**
- **High branch frequency from >4G down to below 2G**
  - Program linkage
  - Function Trace
- **The z196 is not thrilled with these large jumps**
  - Branch prediction affected
- **z/TPF now permits 64-bit CRPAs to reside at 2G to 4G**
  - Lab measurements showed large gains with artificially intense driver
- **Customer measurement/inference showing 6 to 11% performance gain**
- **Joint work with lab and STG**
  - EC12 much better at large jumps
  - APAR still has some value even on the EC12

Existence time calculations

- **TPF Internal Existence time defined as**
  - Message ECB input time into TPF to output time from TPF
- **Little's Law time honored method to calculate response time**
  - Using msg rate and active ECBs
- **Things would be simpler if the message came in, created no additional ECBs, waited until all work was completed, response sent and then exited**
  - This doesn't happen often enough
- **The major source of divergence is the listener/demon ECB**
- **For increased accuracy and partitioning by message group need to use Owners introduced in z/TPF**
  - Use Little's Law on each partition

Example with different estimators

- **Data**
  - Msg/sec=2150
  - Active ECB =680
  - Active IOB=147, IO/sec=89718, IO resp=1.6 mills
  - Pysical IO per msg = 41.9
- **IO wait predictor of msg response time = 41.9 x 1.6 =67 mills**
  - Crude lower bound
- **Little's law with raw ECB count => 680/2150 = 316 mills**
  - upper bound
- **Little's law with adjusted ECB count => 280/2150 = 130 mills**
  - assuming 400 ECBs were listeners
  - TPF data reduction will include a long running ECB count

HyperPAV – Increased Device Throughput

- **Assume for a CU**
  - .5 mills for CU cache hit
  - 10 mills for CU cache miss
  - P(cache hit) = .8; P(cache miss) = .2
  - We have sufficient aliases and host adapter capacity
- **Expected service time = E(s)=.8(.5) +.2(10)=2.4 mills**
- **Thus device rate at full utilization = 1/(.0024) = 416.6 IO/sec**
  - Without HyperPAV- doing one IO at a time
- **Rate to Disk = 416.6(.2) = 83.4 IO/sec**
  - Disk is busy 2 mills out of the E(s)=> 2/2.4 = .83 of total device busy time
  - So if the Disk could get more IO it could do more IO/sec; specifically 100-83.4 = 16.6 IO/sec
    - To send 16.6 IO/sec to the disk we need to send 83.4 IO/sec to the device
- **Now with HyperPAV permitting the CU cache hits to 'MP' with the cache misses we do a total of 416.6 + 83.4 = 500 IO/sec to the device**

HyperPAV Data Collection

- **TPF will report IOBs in use by SSID**
- **Increase or decrease Aliases**
  - Use Little's Law to determine IO response time by SSID

Release performance analysis

- **Increasing trend in mills/msg so only have about 10 days before system change and 10 days after**
- **Assume Tuesday msg rate was 20% higher than other days and hence performance is more critical on Tuesday**
  - On the surface it seems reasonable to do a performance comparisons with ‘Tuesday to Tuesday’
  - Problem is we are giving up too much precision because we only have 2 points before and after
- **Suggest do differences of Mon to Mon, Tues to Tues etc**

Sample Mean

- **Well before the date of the TPF Release or PUT change**
  - compute the sample variance of the daily TPF ITRs
  - With many observations this sample variance will be extremely close to the real underlying variance
  - For Release/PUT change the mean will change but the real variance will not
-

$$\mu_{\bar{x}} = \mu$$
$$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$$

$$\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i$$

Few sample points will have significant variation

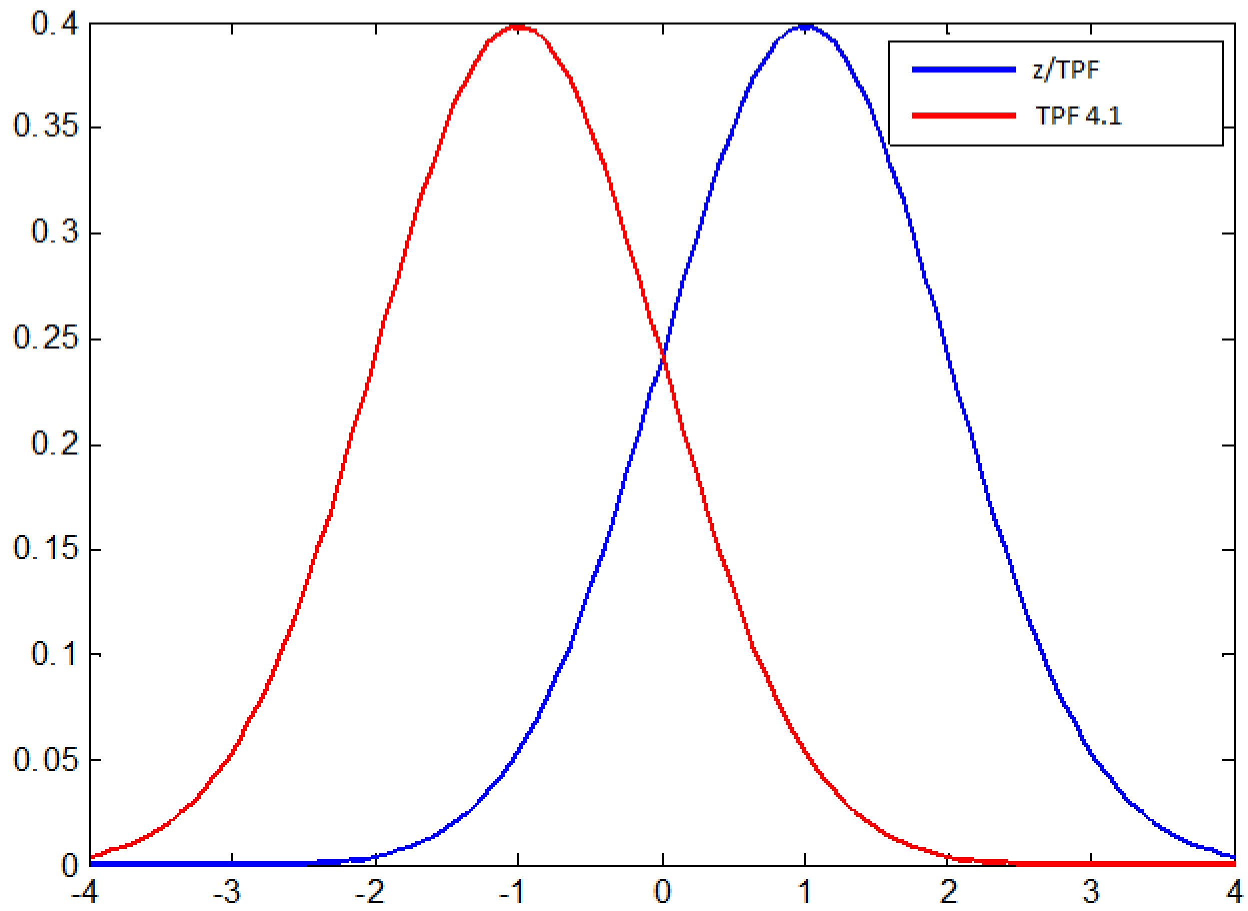

Variance reduction with the sample mean

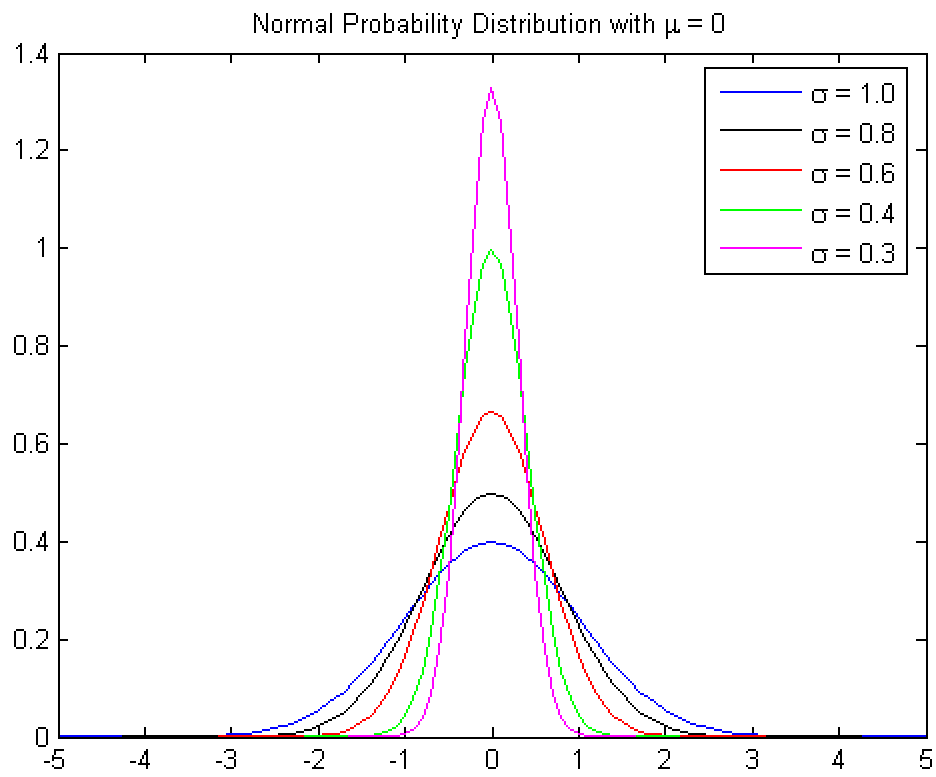

Sample Mean convergence to true mean

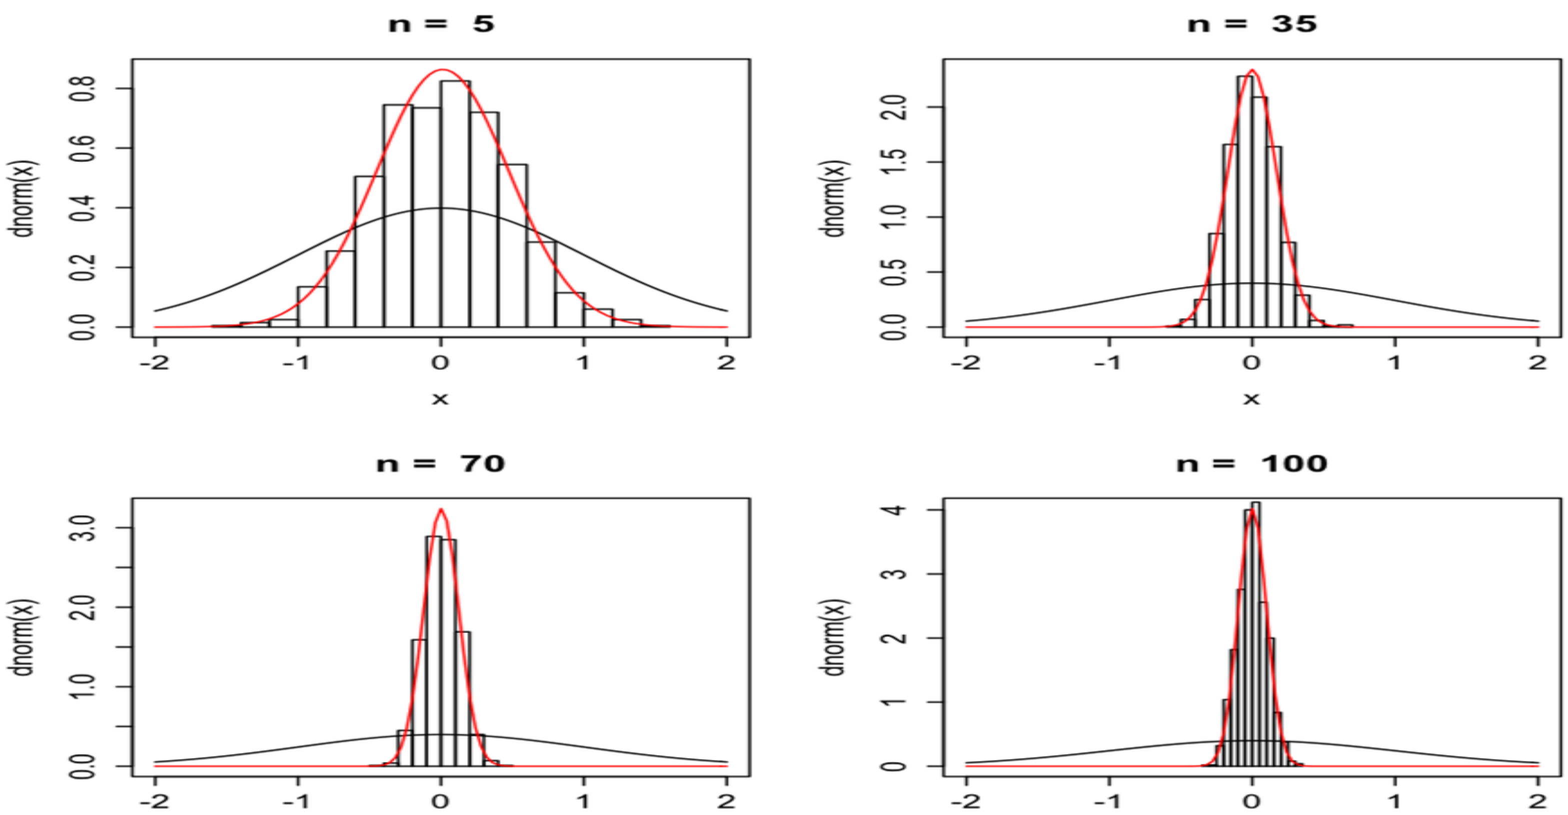

Trademarks

- **IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.**
- ***(Include any special attribution statements as required – see Trademark guidelines on https://w3-03.ibm.com/chq/legal/lis.nsf/lawdoc/5A84050DEC58FE31852576850074BB32?OpenDocument#Developing%20the%20Special%20Non-IBM%20Tr)***

Notes

- **Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.**
- **All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.**
- **This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.**
- **All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.**
- **Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.**
- **Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.**
- **This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.**