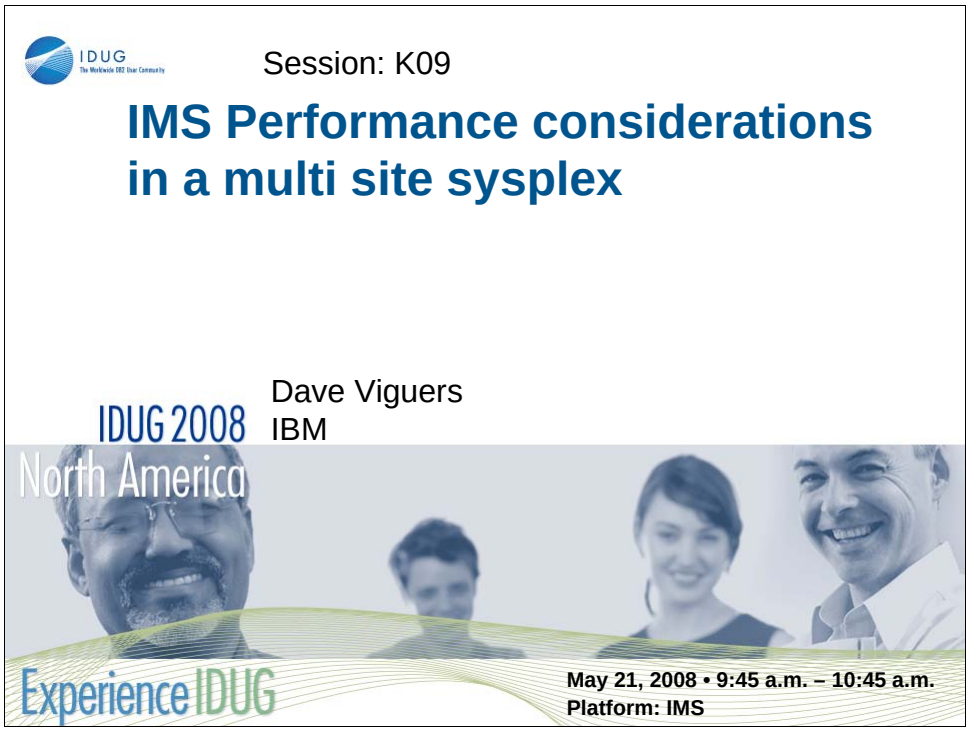

As more and more installations are implementing or thinking about multi site sysplex environments for availability and disaster recovery it is important to understand the performance implications to your IMS transactions and databases. IBM performed some benchmarking of IMS and other subsystems at various distances to better understand these implications and this information will be used to demonstrate the impact and ways it might be minimized.

IMS Performance considerations in a multi site sysplex

- IMS configuration for multi site sysplex
- Coupling Facility response time impact
- DASD response time impact at distance
- Distance impact on IMS transaction response times
- What to tune to minimize the impact of distance

These are the topics to be covered in this session.

IMS Configuration

- Number of possibilities
  - Active site with DASD mirroring
    - Synchronous or Asynchronous
  - Multiple active sites
    - Primary DASD and CF's local to 1 site
    - CF's at both sites
    - Etc.
- For this discussion we will look at
  - Two active sites in a sysplex configuration
  - Two CF's but only 1 being used for IMS
    - Local to 1 IMS but remote to the other
    - DASD local to 1 IMS and remote to the other

There are many variations in possible configurations with a multi site sysplex. The data in subsequent slides is taken from actual measurements done using the configuration shown at the bottom of this slide. Hardware used was a z9 processor for z/OS and CF's with DS8000 DASD. We varied the distance between the two sites (logically of course) for different measurements.

What things might change with distance?

- CPU
- Coupling Facility structure response times
  - Async and sync
- DASD Response time
- Transaction response times
  - DB response times
  - TM (SMQ) response times
- Dependent region occupancy

Here are some of the things which may be impacted as distance is brought into the picture. Each of these will be further explained in subsequent slides.

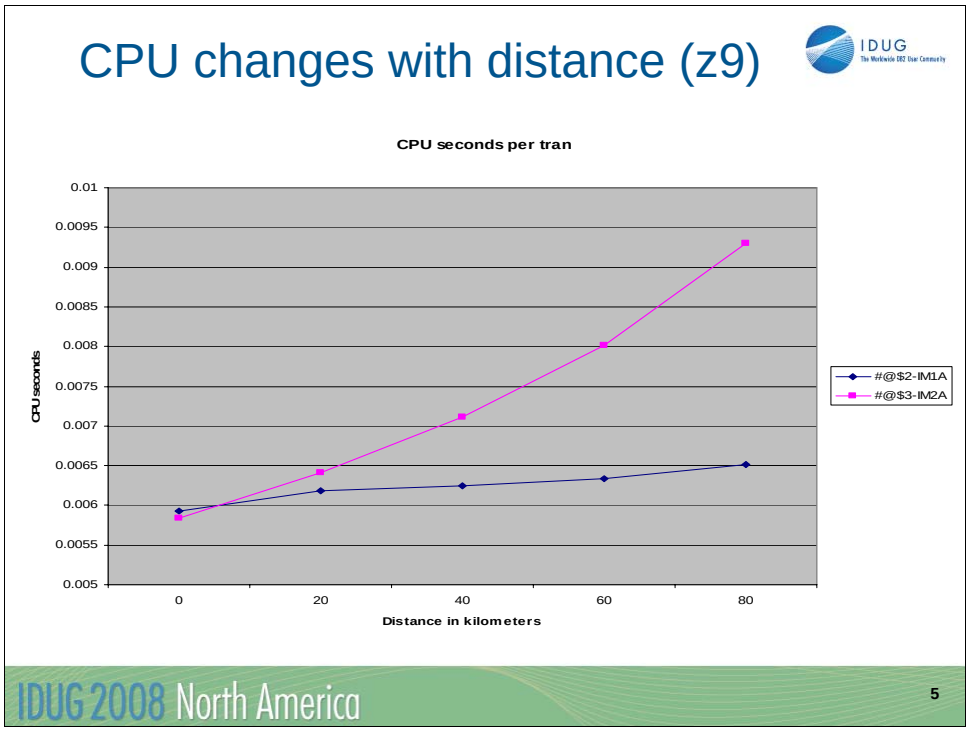

As you can see the total CPU on the local system does not change very much but there is a definite increase on the remote system. Looking at the address spaces in further detail shows that most of the increase is in the XCF address space due to more CF requests being handled asynchronously as the distance to the CF increased.

IMS Coupling Facility Structures

- IRLM Lock structure
- VSAM and OSAM cache structures
- FP shared VSO cache structures
- Shared Queues Structures
  - Full Function Queue structure
  - Full Function MVS Logger structure
  - Fast Path Queue structure
  - Fast Path MVS Logger structure
- The following CF numbers are from 2094 (z9) machines

Here are the types of structures used in an IMS environment. We will see that the system remote from the CF will always show an increase in response time and depending on the type of structure and how it is used there may be an impact on the local system as well. As mentioned earlier all the measurements were done with a z9 for both z/OS and the CF.

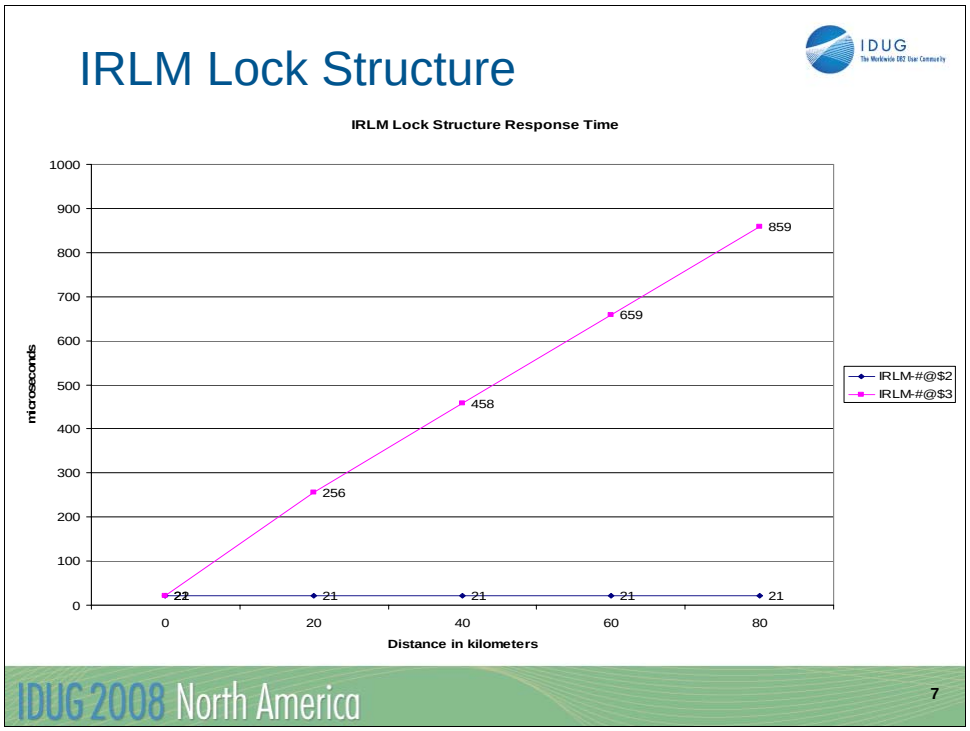

For the lock structure these is no affect on the system local to the CF. The response time from the remote system went up by about 200 microseconds for every 20 kilometers.

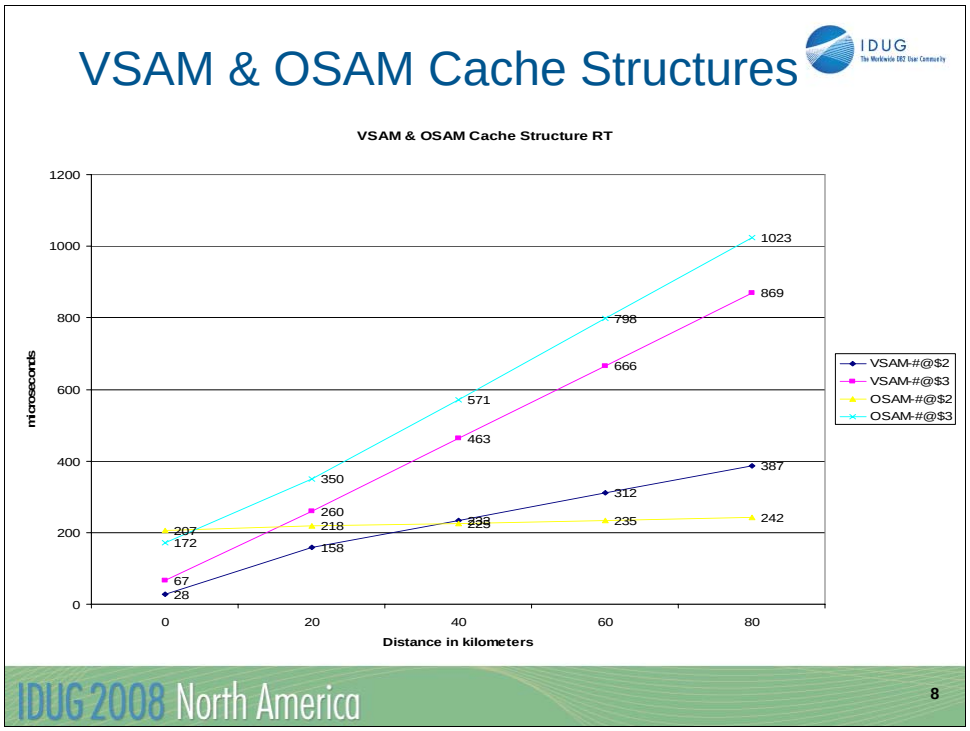

Read requests on the local system see little impact however when IMS updates a buffer in the local buffer pool it must tell the CF to invalidate buffers in any other systems containing the same block of data. This request must be done synchronous and thus the response time is impacted even on the system local to the CF.

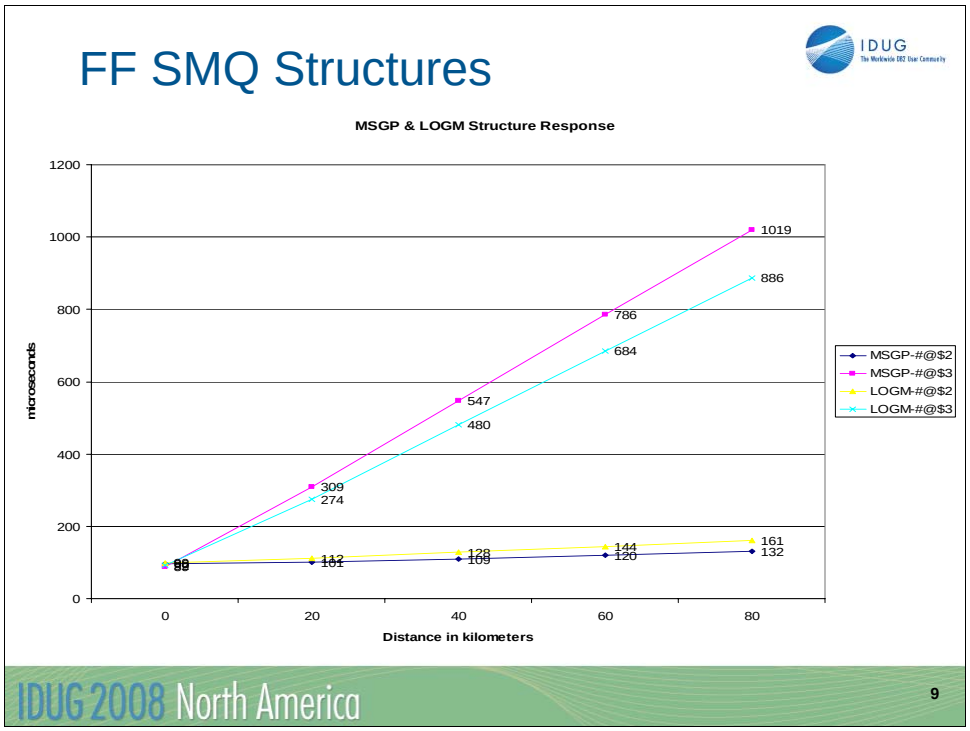

As you saw with the cache structures the local system response times were affected by cross invalidate function. In the case the shared queues list structure it is the list transition notify which will impact the local response time. This happens when a queue in the structure transitions from empty to non empty in which case the other systems in the sysplex must all be notified.

Note: There is no separate graph for the fast path shared EMH structure but the times are very similar.

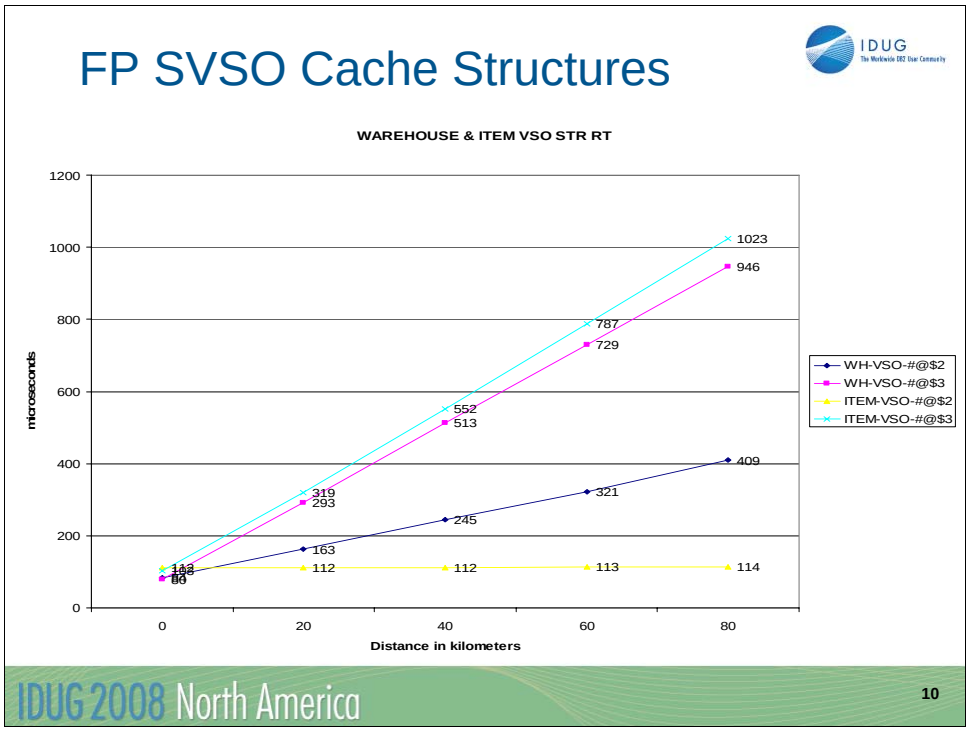

The ITEM structure shows no effect from distance on the local system. This is because it is read only and there are no cross invalidates. The Warehouse structure on the other hand does have updates thus there is an increase in response time even on the system local to the CF.

DASD response time

- WADS
  - Most critical IMS data set
    - Multiple segment chaining can compensate
- OLDS
  - Normal BSAM chaining and buffering can typically compensate for slightly longer r/t
- Database data sets
  - Impact will depend on specific applications
- MVS Logger staging data sets
  - Can be critical if used

Any or all of these data sets can impact response time of a transaction or IMS in general. Even though the WADS are typically the most critical, that is not always the case.

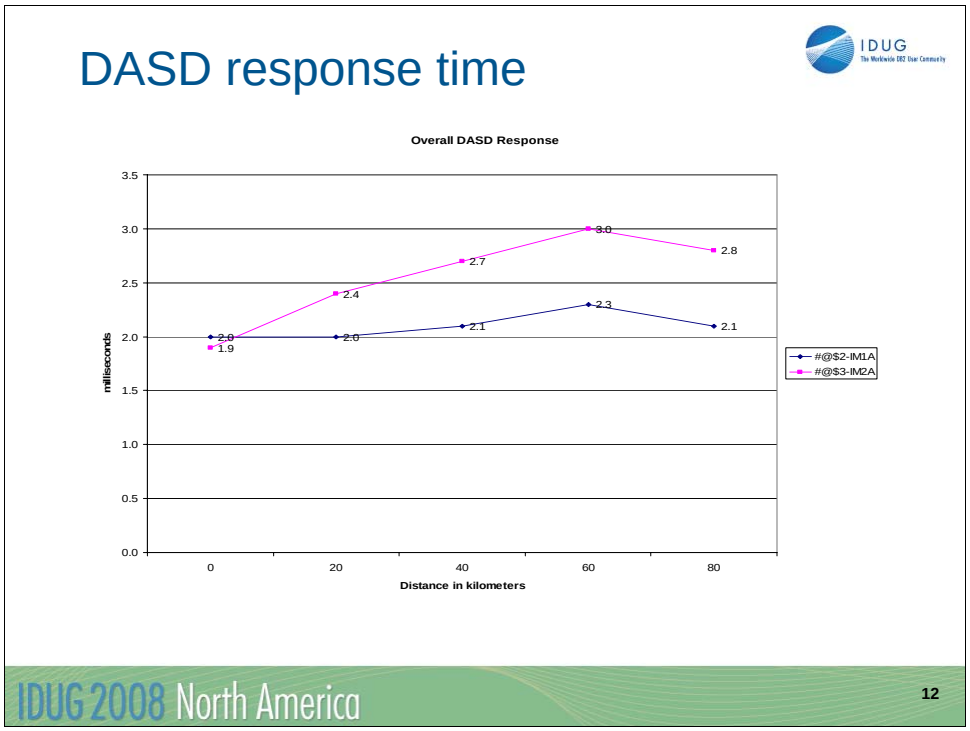

The DASD PEND times increased consistently at about .2ms for each 20 kilometers. The connect time accounts for the rest of the difference. There was an effect most notably with the fast path transactions which will be discussed later which is believed to have caused the slight decrease in time at the 80 km measurement.

Transaction response time

- Full Function
- Fast Path

Next we will look at the affect on transaction response times for both full function and fast path transactions.

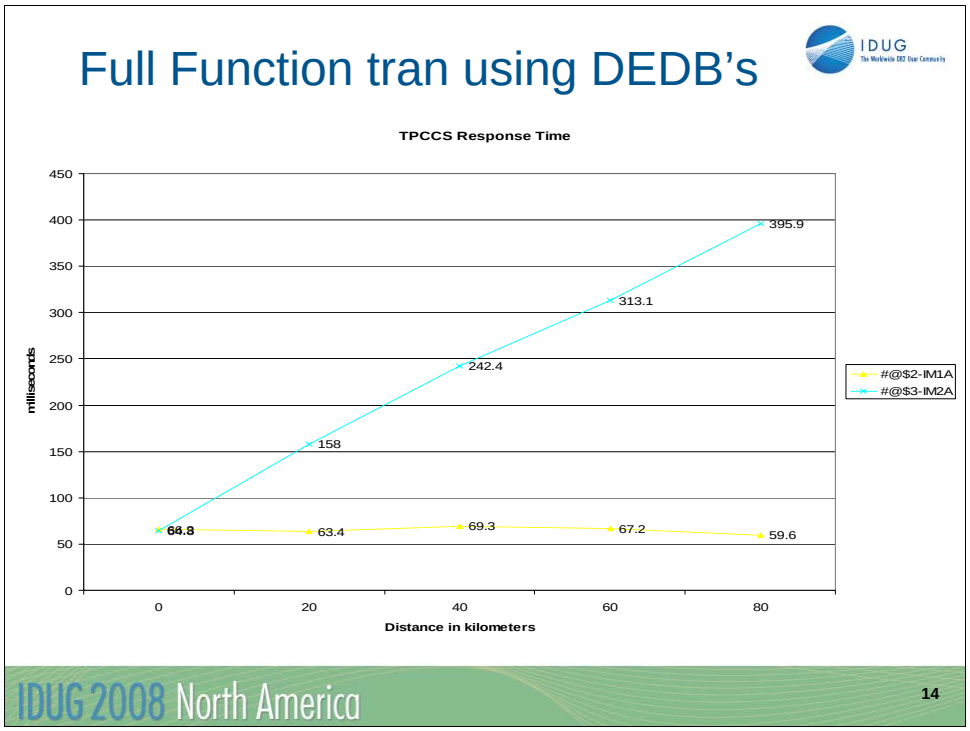

This transaction does 200 GU calls each to 2 different DEDB's. Since there are only reads the local system sees no increase in time and the remote system increases linearly with distance.

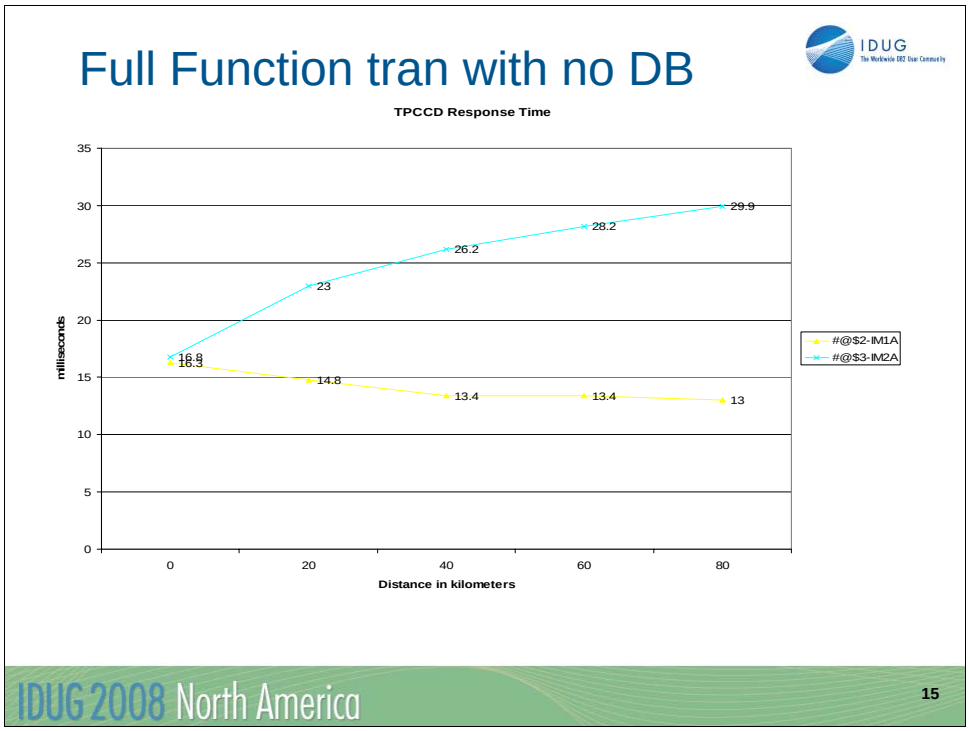

This transaction does only 1 GU to the message queue and 2 ISRT calls for output messages so this transaction is affected only by the SMQ structures. If you think back to the structure response times you will remember they were linear but here we can see that the transactions processed locally actually had better response times as the other systems distance increased. The reason was not obvious but upon looking deeper what we found were that there was an overall increase in total transactions running on the local system. As a result there were more CHKW requests which resulted in shorter WTWT times which resulted in this transaction running faster..

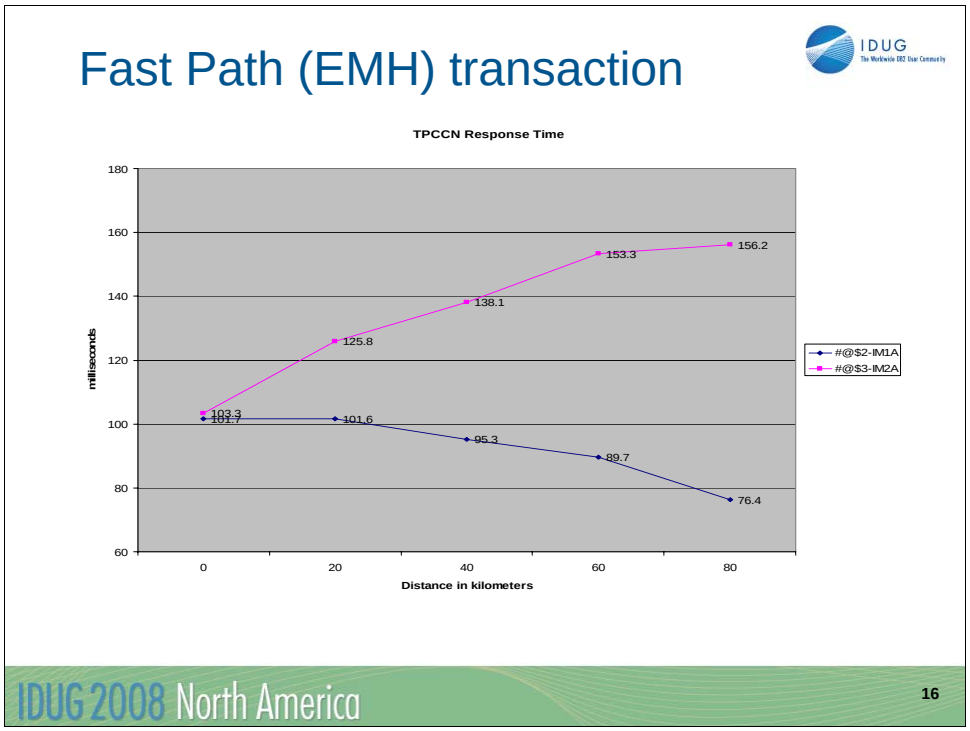

This transaction does 43 calls to DEDB's and 3 to full function. The response time for the local system was actually considerably better at the longer distance. This was somewhat due to the WTWT affect as was shown previously but was mostly due to the fact that we used an artificial driver (TPNS). At shorter distances there was something called the 'wave' effect that smoothened out with distance causing more transactions to be processed locally without ever accessing the CF structure. Without this effect we expect the local line would be fairly flat and the remote line a bit more linear.

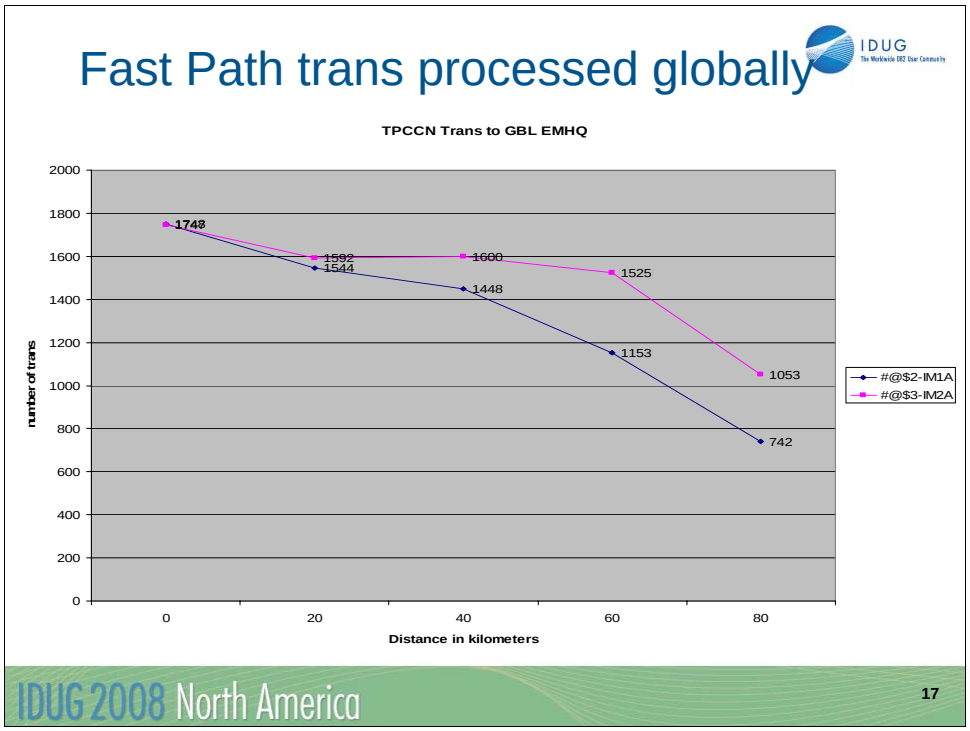

This chart helps explain the results seen in the previous response time chart for this transaction.

Tuning?

- Tuning for distance is not much different than anything else
  - Minimize SMQ CF access with proper SHMSGSZ, LGMSGSZ, & QBUFSZ
  - Try to minimize I/O to critical DB's with buffer tuning
  - Maybe adjust number of regions
    - Depending on current utilization
  - And of course use fastest links possible

Summary

- **Key observations**
- CPU time increases as the distance increases. Mostly XCF.
- DASD response times increase at about 300 microsecs/20km
- The remote lock structure times increased approx. 200 usmicroseconds/20km
- The local system showed no impact for the lock str.
- CF structure access times for the OSAM, VSAM, and VSO cache strutures increase at about 200 microseconds for each 20 kilometers
- For the system local to the CF the times will increase based on the amount of cross invalidate processing.
- The SMQ related structures also showed about 200 microseconds of increase in response time for each 20 kilometers.
- The local system will be impacted by queue transition notifications.
- IMS will not always show a linear increase in r/t as distance increases.
- Don't spend too much time trying to deal with individual transactions unless you have certain high volume transactions which make up a very large percentage of the work.
- In most cases just understanding the overall rate of CF accesses to the IMS related structures should be enough to estimate the impact of distance on your workload.

Session K09

IMS Performance considerations in a multi site sysplex

Dave Viguers

IBM

dviguers@us.ibm.com