Spirent Communications

Inspired Innovation

White Paper

Measuring Jitter Accurately

May 2007

P/N 79-001982 Rev.A 0507

Spirent Communications, Inc.
1325 Borregas Avenue
Sunnyvale, CA 94089 USA

Email: sales-spirent@spirent.com
Web: http://www.spirent.com

Americas
T: +1 800.SPIRENT
+818 676.2683

Europe, Middle East, Africa
T: +33 1 6137.2250

Asia Pacific
T: +852 2511.3822

Measuring Jitter Accurately

Contents

Why Measure Jitter?

Understanding latency characteristics in networks and devices has become critical with greater use of delay sensitive voice and video traffic over IP and Ethernet networks. Two key statistics should be measured when characterizing the temporal performance of a network: latency and jitter. Too much latency renders interactive applications such as voice and two-way video unusable, as the typical person will not tolerate excessive delays in conversation. Similarly, excessive jitter makes the service unusable by negatively impacting service quality.

Jitter is the change in latency from packet to packet. Applications and end-user devices are designed to tolerate a certain amount of jitter. This is achieved by buffering the data flow and designing processing algorithms to compensate for small changes in latency occurring from packet to packet. Excessive jitter, on the other hand, could cause the buffers to overflow, underflow or even cause the algorithm to break down. This creates problems such as dropouts in an audio stream or a choppy video display.

Depending on the application, the tolerable amount of jitter will vary, usually less than 50 ms for most Triple Play services. For example, a good service would have a jitter of 20 ms or less. End-to-end network jitter (or service jitter) can be introduced in a variety of different ways such as packets (in a flow) taking different routes as a result of network congestion or link failure.

The more important cause of jitter is the type introduced by network devices. Buffering, queuing, and switching architectures of any network device inherently have jitter. The jitter varies with traffic characteristics (packet burst distribution, packet length, traffic priority), traffic load, number of users, device load, etc. As traffic traverses a network, jitter is compounded by each device through which the traffic passes.

When designing a network or a network service, an accurate measurement of jitter should be available along with the ability to quantify the jitter amount introduced by network components such as routers and switches. High performance devices (high performance = low jitter) in the network reduce cumulative jitter and therefore provide high quality of service to users.

Defining Jitter

RFC 4689 defines jitter as the absolute value of the difference between the Forwarding Delay of two consecutive received packets belonging to the same stream.

For example, when two packets (packets A and B) are sent through a network:

Packet A takes 15 ms to traverse the network.

Packet B takes 18 ms to traverse the network.

The difference in latency between the two packets in the pair is 3 ms.

Jitter = | 15 – 18 | = 3 ms.

Taking into account RFC 4689, calculation of jitter requires the measurement of four parameters:

- Transmit time of the first packet in the pair
- Receive time of the first packet in the pair
- Transmit time of the second packet in the pair
- Receive time of the second packet in the pair

If A is the first packet and B is the second packet, then jitter can be expressed as:

$$|(RxA - TxA) - (RxB - TxB)|$$

Jitter is always expressed as a positive number, so the absolute value of the difference is used. Average jitter is defined as the average value of the jitter of consecutive packet pairs. If the latency is constant and each packet experiences the same delay, then the jitter would be 0 (since the difference in latency from packet to packet would not change).

Figure 1 graphically shows how jitter is calculated. Two successive packets A and B are transmitted at times TxA and TxB. Packet A takes LA seconds to get through the network and packet B takes LB seconds. The jitter of this packet pair is the absolute value of LA minus LB.

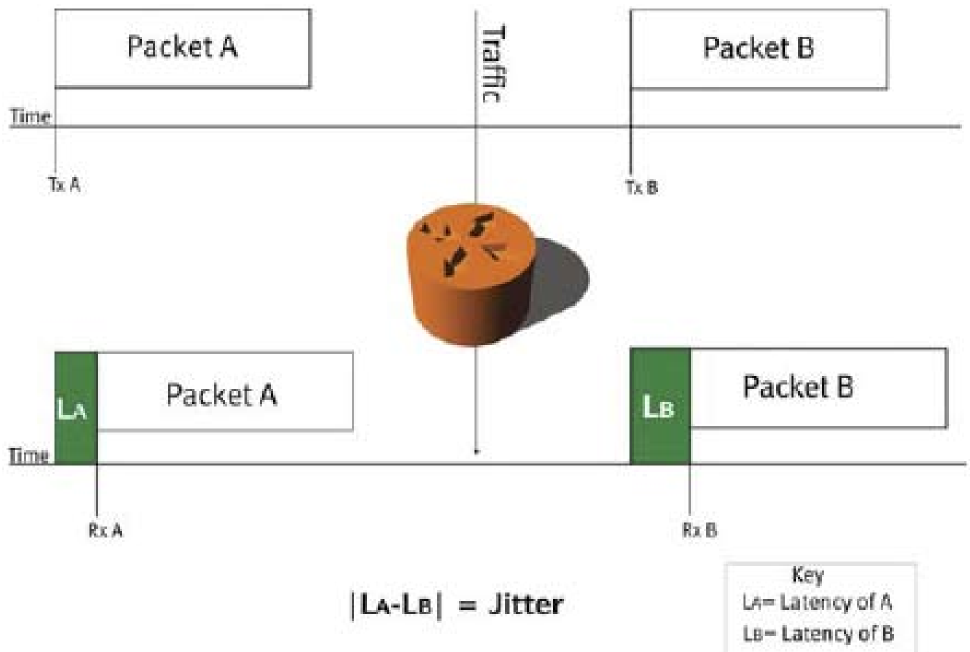

Figure 1

Three Ways to Measure Jitter

Three common methods of measuring jitter are inter-arrival time method, capture and post-process method and the true real-time jitter measurement method.

Hardware required for jitter measurement varies with the above methods, with the data rate to be analyzed and with the desired measurement accuracy. Slow traffic can be analyzed using a PC and the capture method, but that only provides coarse accuracy. The best resolution that can be achieved with a PC is about 1 ms.

In contrast, true real-time jitter measurement at high resolution requires specialized hardware that accurately processes the data in real time at line rates up to 10Gb/s. High performance test equipment today achieves sub-100 ns accuracy.

This paper will examine the advantages and limitations of the three measurement methods. Other methods are not considered because they are insufficient for laboratory test environments. For example, jitter can be approximated as the difference between the maximum packet latency and minimum packet latency over a given period of time. However, this method fails to measure packet pair latency. Moreover, the results can be corrupted by macro changes in latency. For instance, the latency through a device could steadily increase from 20 ms to 200 ms over the period of the test. In this case, jitter would be calculated at 180 ms which would be far higher than the actual packet to packet jitter.

Typical Traffic Scenarios
• Send traffic at a constant rate (10%, 50%, 100%) using fixed length packets • Send traffic at a constant rate using varying length packets • Send traffic at a varying rate using varying length packets (realistic bursty traffic) • Send traffic-pair configuration without congestion

Table 1

A typical test plan for a network or network device would include a measurement for jitter performance under various traffic scenarios. Table 1 lists a few scenarios. The jitter measurement method should be evaluated against test scenario and traffic requirements.

Inter-arrival Histogram Method

The inter-arrival method is a popular way to measure jitter. It uses the trick of transmitting the packets at a known constant interval. By using this trick, two of the four needed parameters are pre-determined. Since packets are transmitted at a known fixed interval, only the inter-arrival time of the received packets is to be measured. The difference in the inter-arrival time between packets is the packet-to-packet jitter. Inter-arrival values are measured over a period of time and displayed in a histogram.

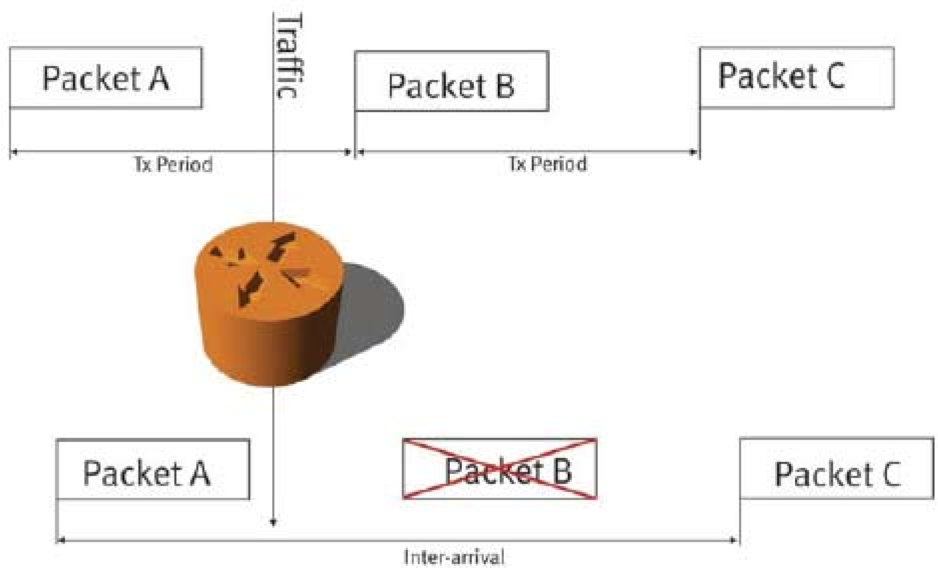

Figure 2

The inter-arrival method has a critical limitation and a few accuracy flaws. The critical limitation is packets must be sent at equal intervals. This restricts the measurement to only constant periodic traffic with fixed packet intervals. Depending on the complexity of the generating hardware, a further restriction requiring fixed packet sizes could also exist. If the hardware can vary packet size but maintain an exact packet-to-packet interval, then varying packet sizes can be used. Because packets must be sent in perfectly equal intervals, it becomes impossible to measure jitter on traffic with a varying rate (bursty traffic).

A key accuracy flaw of the inter-arrival histogram method occurs when a packet is lost (dropped or corrupted). The inter-arrival time between the two packets before and after the dropped packet will be large and will corrupt the inter-arrival histogram. To eliminate corruption of results, inter-arrival measurements should be discarded when packet loss occurs. However, only the most advanced test equipment is capable of discarding the dropped packet inter-arrival data.

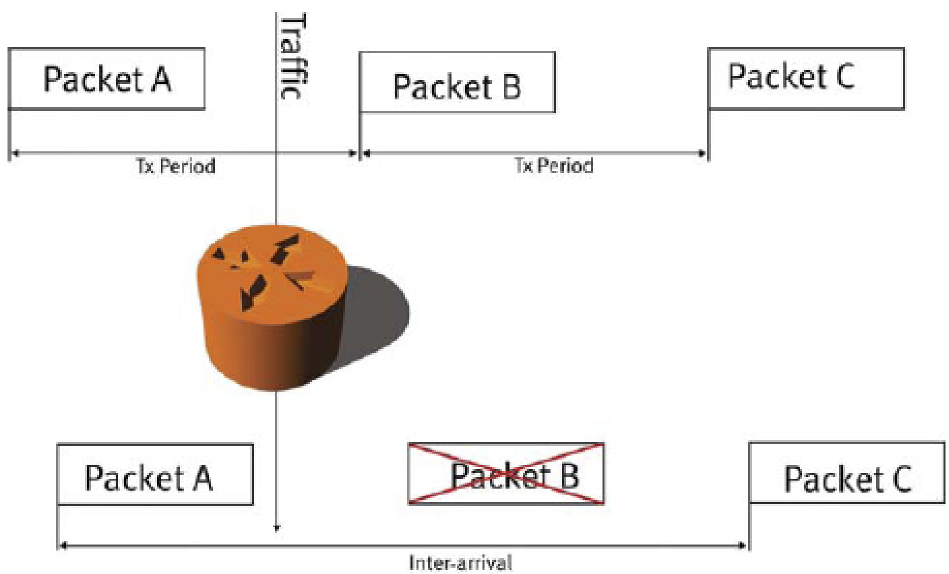

Figure 3. Packet B has been dropped (indicated by the red X) and did not arrive at the destination. The inter-arrival time between Packet A and the next packet received (Packet C) has been calculated incorrectly since this method does not properly account for the lost packet (Packet B). The inter-arrival method will indicate an erroneously high jitter value due to these dropped packets.

A related accuracy flaw in the inter-arrival histogram method is that the method fails to take into account packets arriving out of order. Packets arriving in a different order from the order in which they were sent also corrupts the measurement.

Capture Method

A second common way to measure jitter is to capture all packets and then process the data offline. Most test equipment puts a signature in the sent packets and thus the capture file contains all the needed information (timestamps in the packets indicating Tx times and capture buffer hardware indicating Rx times). Test signatures also include packet sequencing information, making it possible to compensate for lost or out-of-sequence packets when using this method.

The critical limitation of the capture method is finite space in the buffer. The buffer can be filled up very quickly if data is sent at high speed. Typical test plans dictate the need to measure jitter over a much longer period of time than is possible with the largest capture buffers on most of the current test equipment. Another limitation of the capture method is the lack of real-time, cause-and-effect analysis. Debugging and analysis time is greatly reduced if the engineer can change a traffic load or device configuration parameter, and see feedback in the jitter measurement. This real-time operation is not possible using the capture method.

True Jitter Measurement

To provide a set of industry standard definitions, Metro Ethernet Forum (MEF) released the MEF 10 specification in 2004. This specification contains a section defining the proper way to measure jitter while taking into account lost or corrupt packets. Figure 4 is a flow chart illustrating how Spirent TestCenter 2.0 implements the MEF 10 jitter measurement definition.

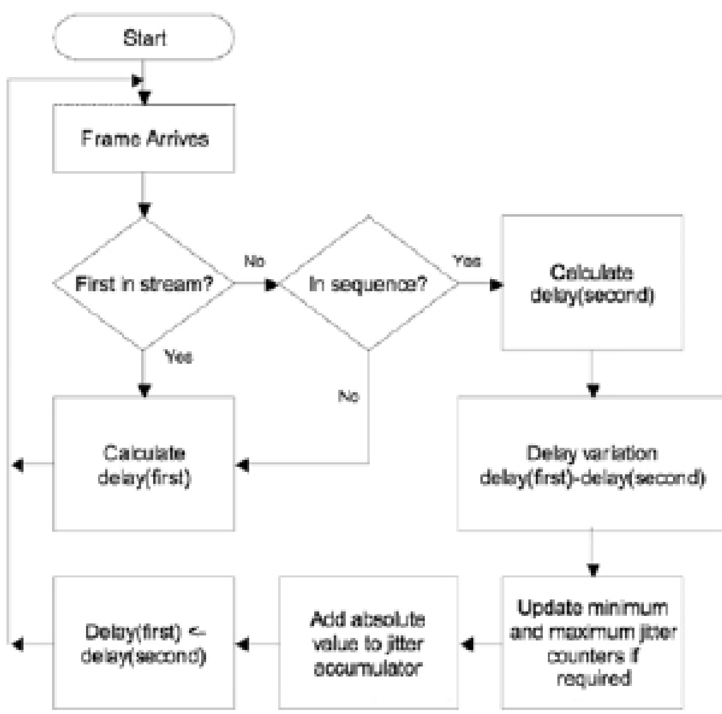

Figure 4

If the packet is the first received in the stream, then the packet transfer delay (latency) is calculated and stored. If a received packet is not the first packet in the stream then a check needs to be performed to make sure the packet is in the correct sequence. If the packet is not in sequence, latency results are discarded and this packet is treated as the "new" first packet in the stream. This stops measurement corruption caused by lost or out-of-sequence packets. If the received packet is not the first packet and is in sequence, then the delay is calculated and stored. Next, the delay variation (jitter) is calculated by taking the difference of the delay of the current packet and the delay of the previous packet. Maximum, minimum, and accumulated jitter values are updated and stored. Finally, delay of the current packet is saved (to be used as the previous packet delay when the next packet arrives). This algorithm runs in hardware at full-line rate up to 10Gb/s.

The main advantages of the true real-time jitter measurement are no dependence on packets needing to be sent at a known interval; and, the method can measure jitter on variable rate (bursty) traffic. Also, this method does not restrict test duration. This is because the calculation occurs in real time as packets are received with no need of packet capture. Finally, this method compensates for lost and out-of-sequence packets while producing results in real-time for instant feedback even when varying traffic or device parameters.

Other advantages of true real-time jitter measurement include complex analysis views such as jitter charts or jitter histograms. These views produce far more revealing pictures than other measurement methods. The views substantially reduce test and analysis time. For example, the typical inter-arrival method produces a histogram of inter-arrival times showing how many packets were received in each inter-arrival bucket. When using only the inter-arrival histogram, it is not possible to determine when, where or why abnormal amounts of jitter occurred. All that can be determined is the approximate maximum, minimum and average jitter.

Looking at Histogram 1, the jitter (width of the histogram) is measured to be about 250 ms. The corresponding line graph of the measured jitter is shown in Graph 1.

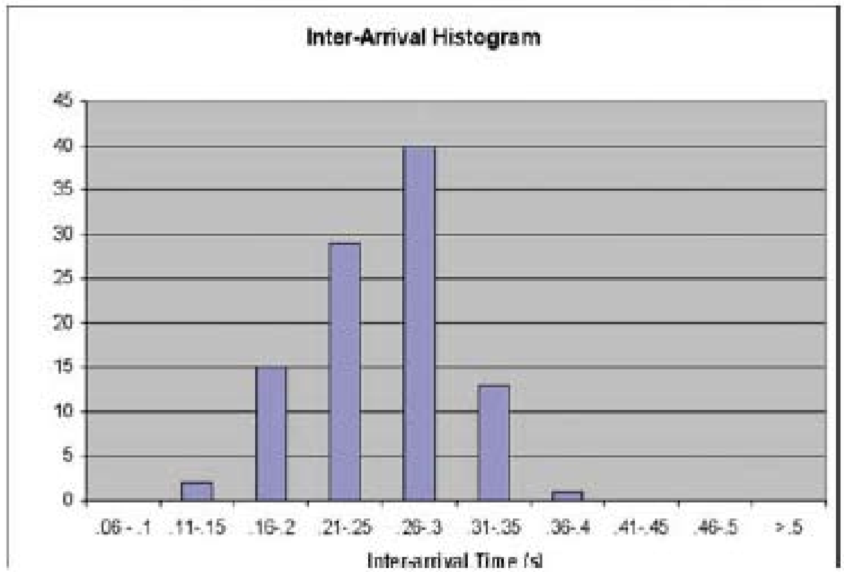

Histogram 1

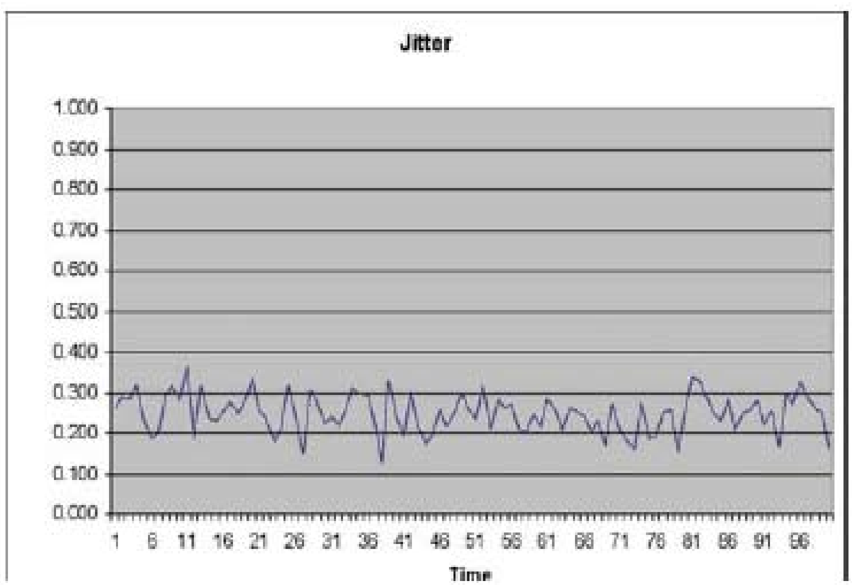

Graph 1

However, Histogram 2 below does not provide such a clear picture. Roughly one-fifth of packets have a much higher inter-arrival time. From the second histogram, it is not possible to tell how the jitter was distributed over time. Jitter may have occurred as a single burst (Graph 2A) or it could have been distributed as several bursts of higher jitter (Graph 2B).

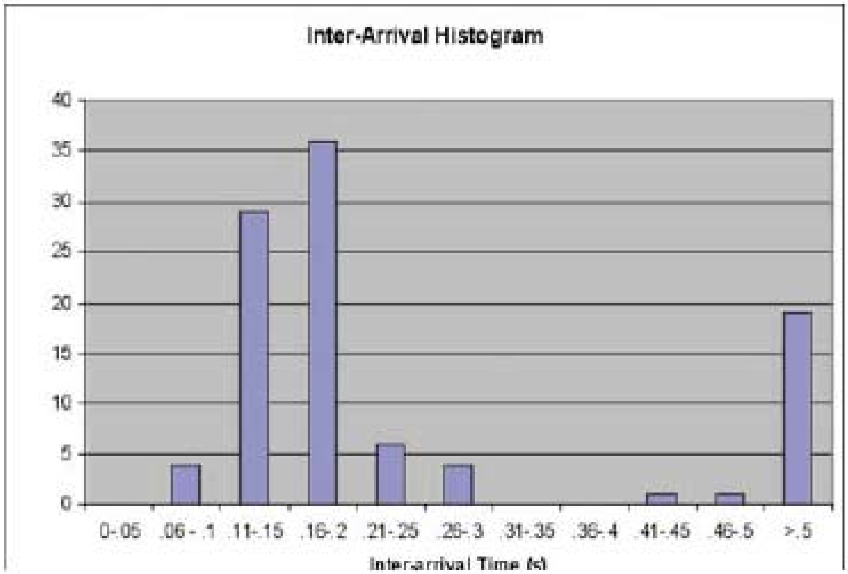

Histogram 2

Both Graph 2A and Graph 2B produce exactly the same histogram shown in Histogram 2. This complicates the problem analysis and creates a difficult task for the test engineer. If the user sees a live graph of jitter vs. time, the problem is discovered faster and analysis proceeds more quickly.

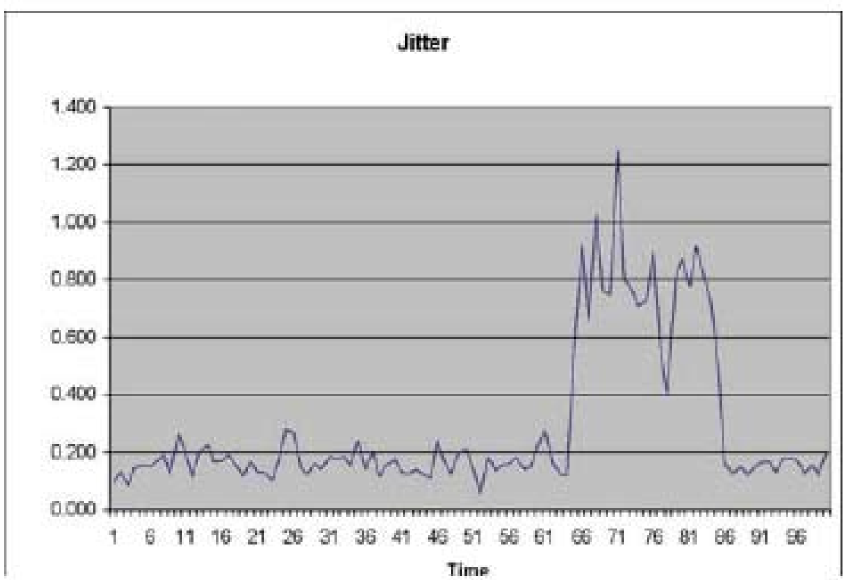

Graph 2A

Also notice that if histogram bins are not properly set, the jitter peak cannot be measured properly. The maximum inter-arrival bucket in the example was set to .46 to .5 seconds. Inter-arrival information above .5 seconds was lost because the maximum inter-arrival time was exceeded during the test. This would require the engineer to re-run the test which may be impossible if the event is not easily reproduced. Valuable information will be lost.

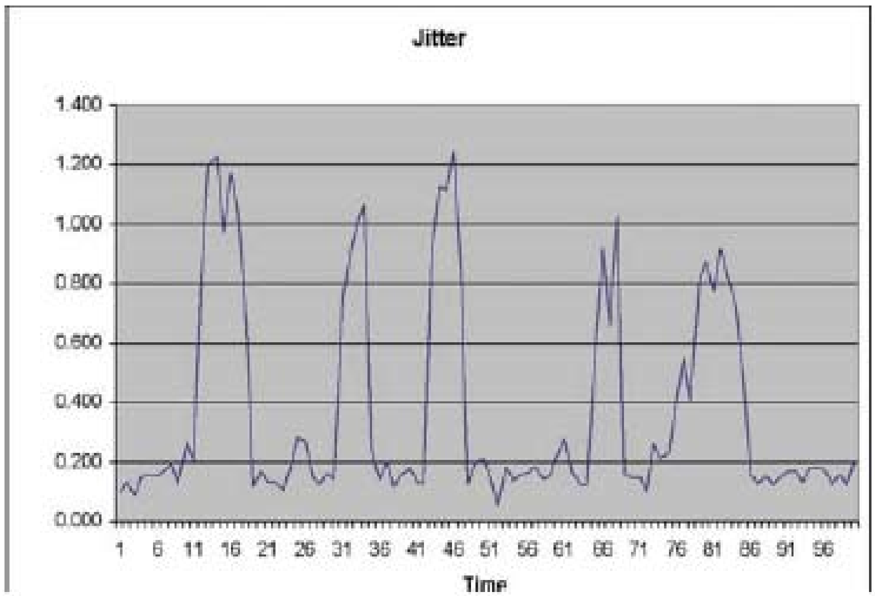

Graph 2B

Why Real-Time Jitter Measurement is Best

By analyzing the following Table 2 which summarizes the key requirements for measuring jitter, it is evident that true real-time jitter measurement is superior to the other two methods. Real-time jitter measurement provides test scenario flexibility, accurate results and real-time analysis capability. With capability and performance advancements in state-of-the-art test equipment, this highly accurate and powerful jitter measurement method is now available to test engineers. By having an accurate and clear picture of jitter performance, the test engineer can better understand jitter characteristics in the network or network device.

Key Requirements for Measuring Jitter

	True Real-Time Jitter	Inter-arrival Histogram	Capture
Accommodates all required traffic scenarios	Yes	No	Yes
Able to continuously measure jitter performance over long test periods	Yes	Yes	No
Properly calculates jitter when packet loss occurs	Yes	No	Yes
Provides detailed real-time jitter results (graphs and histograms) and allows efficient problem analysis)	Yes	No	No
Able to measure jitter as defined in MEF 10	Yes	No	No

Table 2

Inspired Innovation