BANDWIDTH DRIVEN TELEMETRY

Charles H. Jones and Lee S. Gardner
412th TW/TS
Edwards AFB, CA 93524

ABSTRACT

> "The radio frequency spectrum is a limited natural resource; therefore, efficient use of available spectrum is mandatory."
>
> IRIG Standard 106-96 [4]

As the availability of the frequency spectrum decreases and demands for bandwidth from users increases, the telemetry community will have to find ways to use spectrum efficiently. This paper is an overview of the major areas of research that promise potential increases in the efficient use of the telemetry spectrum. The discussion is summarized in a matrix that compares potential gains with overall costs for each research area using relative values of high, medium, and low.

INTRODUCTION

The telemetry community is recognizing that available frequency spectrum is a scarce resource. Two major contributions to this growing awareness are: an increase in demand; and recent reductions in allocation of frequencies due to government auctions. As an illustration of the approaching crunch, consider a scenario that is being discussed within the Flight Test Community – that of a 'network in the sky.' This concept is simply the extension of networking technology to include aircraft as nodes in the network – connected by telemetry links. Now consider the abundance of 100 Mb per second networks and the fact that these are rapidly becoming too slow for many applications. Finally, consider that, with current technologies and operational constraints, *there does not exist* 100 Mb of bandwidth for use in flight test telemetry. (Roughly speaking, 210 MHz at about 3 hertz per bit gives about 70 Mb.) Admittedly, ground-based networks are different than this 'network in the sky' concept, but even if telemetry bandwidth requirements are, say, 5 to 10 years behind ground-based networks, then 100 Mb telemetry requirements are only 5 to 10 years away.

Some other concerns about this crunch are bit error rates and an increase in the number of channels in use simultaneously. Reducing the bit error rate is accomplished either by increasing power – usually by increasing spectrum used for the same bandwidth – or introducing error correction techniques – which add bits. As an example of the increased channel usage, an Advanced Range Telemetry (ARTM) survey [1][2] found up to 28 channels being used simultaneously on the same test range. Before going into specific technologies that show promise to alleviate the coming crunch, let us overview some aspects of how we currently use spectrum.

Basically, there are two modes of communicating information: synchronous and asynchronous. In the flight test telemetering domain, these roughly correspond to methods like pulse code modulation (PCM), which is synchronous in nature, and packetized telemetry, an asynchronous method. The PCM typically consists of a continuous stream of bits encoded as analog signals in codes like NRZ-L, or BI-Φ. One knows the information content of the stream by counting bits from a starting, or synchronization, bit pattern which reoccurs cyclically. Packetized telemetry, on the other hand, may also be encoded as analog bits using the same codes, but is sent in bursts, as needed, rather than in a continuous stream of bits containing a cyclic list of data items. Packets contain sufficient extra bits to identify the specific information they carry plus the associated data bits.

There is currently talk of the basic PCM telemetry stream going away to be replaced by packetized, networked telemetry. This raises a fundamental question of whether data flow will be an ‘open pipe’ (i.e., a continuous, synchronized, cyclic flow of bits) or packetized (i.e., information collected, put into little packets and sent in bursts). The authors contend that there are in fact different applications requiring these very different communication methods.

If telemetry is being used to monitor safety of flight, then there is some question of whether any delay in receipt of the information from the test vehicle is acceptable. Because the packetizing process typically implies buffering the incoming information, packetized telemetry experiences delays and may be inherently unsuited for time-critical applications.

Other applications unsuited to packetized telemetry can be characterized by a requirement for periodic samples. This requirement exists for tests where frequency domain analysis of the test data is required, for example during structural 'flutter' tests, or during flying qualities testing where flight control laws are typically analyzed for frequency response.

In addition to these applications, if there is an increase in demand for data, like there would be in a packetized system during a typical flight test maneuver, the behavior of the telemetering system under 'surge' conditions becomes a concern. As workload increases, the amount of inefficiency inherent in a packetized network structure becomes an issue. Most studies show that contention-based network protocols like ethernet, for example, require about 40 percent of their bandwidth for overhead processing. That is, only about 60 percent of the bits using the available spectrum carry the information critical to the application. Other protocols may be somewhat more efficient, but the very nature of packetization increases overhead so that the efficiency of any network employing packets is asymptotically bound to about 60 percent.

On the other hand, there are many applications where communication with test vehicles does not have the real-time constraints or periodic sampling requirements. For these aperiodic, nonreal-time, nonsafety critical data, packetized telemetry may in fact, afford a viable solution.

For purposes of discussion within the context of this paper, the technologies addressed may apply to either model of telemetry: open pipe, packetized telemetry, or both. Similarly, the technologies discussed may improve spectrum efficiency usage of a single frequency or multiple frequencies.

The rest of this paper overviews various technologies related to increasing spectrum usage. This overview is summarized in a matrix comparing potential improvements to associated costs. Since this paper is an overview, the discussions of the technologies are brief and concentrate on providing justification for the entries in the matrix. The columns in the matrix include amount of research (theory, hardware, and algorithms), amount of re-instrumentation (onboard and ground), and recurring costs.

In order to make these types of judgements about a particular technology or research area, it is first necessary to consider what is meant by the efficient use of the spectrum. A first level metric is simply how many bits per hertz can be transmitted. But, just because a bit is transmitted and received does not mean it carries useful information. The second level metric is how much information can be transmitted per hertz. However, just because a bit is translated into a recognized piece of information, does not mean it is useful to analyzing the data for a particular test. Even further, not all bits carry the same amount of information. For example, if a word of data carries the accumulation of many raw samples (e.g. average, maximum, minimum, or some more complicated derivation) it, in some sense, carries more information than any of the raw data bits. But this leads into a third level metric involving requirements for data. Are the raw samples needed? Or is an average adequate? Finally, since some telemetering applications use more than one channel at a time, efficiency needs to be measured over all channels in use. So, as our measure of bandwidth efficiency, we choose the amount of *required data* that can be sent over the *entire available spectrum*.

REQUIREMENTS ANALYSIS

Exactly what needs to be telemetered? The answer to this question is not easily obtained. Even the most basic question of what measurements are truly required for a given test is not always straight forward and, without extensive analysis, there are probably many nonessential measurands that are added to the requirements list 'just in case.' Similarly, the sample rate for a given measurand is probably not rigorously analyzed to the last hertz. And other, more subtle questions are rarely even considered. What accuracy is needed in a telemetered sample? If the strip chart scale is in 100s, is it necessary to transmit tenths? Or, how much data can be viewed? How many measurands changing at what rate can someone monitor when looking at strip charts or displays in a control room?

All of these questions are related to the basic issue of requirements. Without bandwidth shortages there is not much need to do extensive analysis on exactly what is required in a telemetry stream. Since bandwidth is only starting to be a scarce resource for flight testers, sample rate analysis is not currently extensive; neither is a careful evaluation of precisely what measurements are needed. If this type of requirements analysis is required to be done to ensure efficient spectrum usage, support tools will have to be developed. Some theory exists about how to determine appropriate sample rates, but there are a lot of unknowns requiring research. For example, questions like, "Is there any software that is based on this theory?" or "What theory exits about what is needed to test certain aircraft (or other test vehicle) characteristics?" and others require answers.

As will be seen later in the matrix, the benefits of several of the technologies discussed in this paper can only be realized if better requirements analysis is done. In fact, better requirements analysis may ultimately be where the most benefit can be derived. However, since there has always been more than sufficient spectrum in the past, there has never been a reason for this level of analysis. It is certainly not part of the current mode of operation. Thus, this benefit will only be obtained through a cultural change—always the most difficult change to accomplish.

DYNAMIC MEASURAND SPECIFICATION

Wouldn't it be nice if the set of measurands being telemetered could be changed as needed during the flight? This dynamic specification of measurands takes on at least three aspects: predefined measurand sets either associated with different tests or for viewing different aspects of a given test; on-the-fly measurand requests that are determined during a test; and requests for measurands to be played back from a previous part of the test. If all of these things could be done, test engineers would have a level of flexibility unknown of today; a flexibility that would probably increase the effectiveness of flights by making it much easier for testers to get what they need, when they need it, during the flight.

There are several things that need to be in place to implement this flexibility. An upload telemetry link to be able to tell the test vehicle of the changing measurand set. (This is an excellent example of an asynchronous link.) The ability to automatically generate and switch data cycle map (DCM) formats (see below) in real time must exist. The ability to playback data will probably require solid-state recorders on the aircraft. And, permeating all of this, a method for automated verification of data collection, format validity, and proper coordination between transmitting and receiving systems must be in place. The last is important as much from a user acceptance and confidence point of view, as from a real need to validate a dynamically changing software system.

It might appear that this technology eases the necessity of up front requirements. If you can change what measurands are being sent dynamically, maybe you don't need to be quite as careful about identifying what measurands are needed prior to the flight. However, safety measurands need to be well identified so that they don't get overridden in the haste to see something else. Sample rates still need to be established prior to test. Parameter sets for particular tests still need to be identified carefully so something isn't forgotten by trying to establish them off the cuff. In general, a lack of careful requirements analysis prior to a test will lead to the same inefficiencies currently in place. Finally, the process of determining a new set of measurands in the control room will take time – something that is often critical during an actual test.

DYNAMIC SPECTRUM ALLOCATION

Although real-time monitoring uses an open pipe telemetry, the requirement for information is not necessarily continuous through an entire test and there are other communications going on besides the test monitoring. Thus, dedicating frequencies for the entire planned length of test potentially leaves spectrum unused for long periods of time. It is conceivable that a centralized controller could dynamically allocate the timing, content, frequencies, and power of all test-related transmissions. Potentially, 'dynamic' means fractions of seconds. The obvious analogy is that of a network controller; a central system allocating the spectrum in some systematic manner. However, a normal network is working over a single cable so that the allocation is in some sense one dimensional – only one system can broadcast at a time. In contrast, the telemetry environment is at least three dimensional. Not only can more than one system broadcast at a time using different frequencies, but they can broadcast at different powers and in different directions. This leads to a much more complicated decision process. There are other complications as well. The requirements must be determined dynamically, in real time. Each system participating must be able to communicate its on-demand, changing requirements to the controller. One final complication is that the controller must communicate its allocations to all involved parties.

The use of this technique has potential for a high increase in efficiency since it attempts to avoid any down time in any frequency. The corresponding costs are also high in that neither the algorithms nor the hardware for doing this currently exist, nor does the expertise to analyze test requirements on-the-fly.

DATA CYCLE MAP (DCM) DESIGN

A DCM (often thought of as a PCM matrix) contains the description of the bits in a telemetered bit stream. The DCM design has been shown to be very difficult [3] and a study [2] has shown that DCM designs tend to be inefficient in practice. Better DCM designs lead directly to better use of telemetry spectrum. In order to achieve better designs, it is necessary to overcome the inherent difficulty in the design problem. This requires better theoretical understanding of the underlying structure of DCMs and the development of automated DCM generators. In addition to this theoretical research, bit rate agile hardware must be developed and put in place before potentially more efficient designs can be implemented. That is, hardware is required that will support any bit rate that a DCM designer comes up with based only on test requirements, otherwise significant bandwidth is wasted due to arbitrarily filling-up the DCM bits in order to accommodate a limited choice of frequencies. Projects at Edwards AFB, California, are pursuing theory and automation. Vendors are beginning to produce bit rate agile hardware. Thus, these benefits are likely to be achieved in the near future.

However, it must be remembered that the efficiency of DCMs is partially dependent on the accuracy of the input requirements. Thus, full realization of efficient DCMs will come only with complete requirements analysis.

DIRECTIONAL ANTENNAS

One of the difficulties with using many channels is that of interference. The use of directional transmission can reduce interference by providing a single signal path from transmitter to the receiving antenna. Since radar arrays that control the shape and direction of their beams are fairly common, it seems likely that this technology could be modified for telemetry use. Thus, implementing some modified phased array approach to telemetering may require little more than changing antennas and associated software on test vehicles. However, unless very narrow beams can be maintained, the benefits can only be derived by geographically separating test vehicles and receiving antennas. That is, it is most likely to eliminate interference between test ranges, say China Lake, California, and Edwards AFB, California,, rather than on a single range.

DATA COMPRESSION

There are several types of data compression. The first to come to mind is a dynamic version of a Zip™ like compression. That is, a compression method that only looks at the bits without any knowledge of what they stand for, and compresses the data on the fly. Unfortunately, because of the dynamic nature of the algorithm employed, expected compression ratios are very low. A second type of compression is differential compression. Instead of transmitting all bits in a data word, only transmit the low-order bits that change. Thus, instead of transmitting 10 bits, perhaps only 4 are sent. This has potential, but requires the airborne system to recall the last value sent in order to compute the differentials. It also requires some dynamic way of communicating between airborne and ground systems what DCM bits, in which words, pertain to what samples. Finally, the synchronization lock between transmitter and receiver must be maintained, or a differential update could be lost, throwing off all future values as interpreted in the test control room. A third type of compression uses onboard pre-processing of the data. Instead of transmitting all raw samples collected, process the raw data into useful results and only transmit the results. This method, however, is highly dependent of correct analysis of test requirements and may be quite sensitive to loss of one or more sample inputs due to sensor or instrumentation system failure. It is not clear, for example, how someone monitoring a transmitted average of several inputs might distinguish loss of one of the input signals from a normal drop in the average.

All of these techniques have several difficulties. First, compression of any type takes processing time. This introduces delays into the transmitted data. Second, even with lossless compression techniques, the reconstruction of the original data is dependent on receiving all the bits correctly. Bit error rates in telemetry tend to be on the high side and complete dropouts when no bits are detected at all are not uncommon. While the loss of a single bit in uncompressed data most likely will simply cause a spike for a single measurand, the loss of a single bit in compressed data most likely will cause complete loss of data until the next syncronization point. The solution to this potential data loss difficulty is to do forward error correction. But forward error correction techniques depend on adding bits that could significantly offset any gains from compression.

Again, as with previous techniques, the benefits in bandwidth efficiency from using data compression techniques are highly dependent on effective telemetry requirements analysis. Depending on the requirements, any one or all of these compression methods may be acceptable.

MODULATION ALGORITHMS

When thinking of bandwidth increases, changing the modulation technique is perhaps the first to come to mind. How are bits encoded in the spectrum – PCM, frequency modulation (FM), multiplexed? There has been quite a bit written about the subject. However, this direction of exploration is starting to have a problem similar to the 2010 problem. That is, if current miniaturization of memory storage trends continue, by the year 2010 memory will have to be stored at the rate of 1 bit per molecule, atom, or electron. Similarly, one can ask if it is possible to transmit more than 1 bit per wavelength. Technology isn't quite at this point yet, but it's getting close. It seems that the amount of bandwidth to be gained through better modulation suffers diminishing returns in terms of the cost of research necessary to make those gains. Some potential exceptions to this are spread spectrum techniques since they tradeoff power in a single frequency for low power in many frequencies. Unfortunately, in their current incarnation, they have limited range and thus, limited value in flight test. Although, these techniques may be useful with test vehicles that do not travel large distances, or for telemetry between nonmoving platforms.

MATRIX SUMMARY

The entries in this matrix are based on the authors' experience and discussions with colleagues. They are certainly up for debate and, to a great extent, the purpose of this paper is to instigate this debate. The entries are in terms relative to the other technologies for the given category.

The potential increase in bandwidth is based on increasing the amount of *required* data over the *entire* available spectrum. As such, techniques that concentrate on requirements and provide benefits over more than one channel have been given higher entries than those that improve the efficiency of a single channel.

The research column tries to capture a combination of the difficulty of the problem and how much theory exists regarding the problem. The lower the amount of current theoretical underpinnings the higher the risk for a technique.

The re-instrumentation column tries to capture whether systems (hardware and software) on the test vehicle or ground station will have to be replaced, developed, or whether they exist and can be purchased.

The combination of research and re-instrumentation is representative of one time costs. The recurring costs are mostly affected by their dependency on requirements analysis. That is, once a system is developed and put in place, the recurring costs are normal maintenance costs. However, accurate requirements analysis is not likely to be automated in the near future so that if the efficiency increase is dependent on accurate requirements, the recurring costs include the effort for this analysis.

Item	Bandwidth Increase Potential	Research	Re-Instrument	Recurring Costs
Requirements Analysis	High	Med	Low	High
Dynamic Measurand Specification	High	Low	Med	Med
Dynamic Spectrum Allocation	High	High	High	Med
DCM Design	Med	Med	Low	Low
Directional Antennae	Med	Med	High	Low
Compression	Low	Med	Low	Low
Modulation	Low	Med	Med	Low

CONCLUSIONS

Whether it is called a cultural change, or a paradigm shift, or some other phrase because you're tired of these, dramatic change is required. Efficient use of bandwidth has not been a requirement—but it is now. This requirement affects every aspect of the flight test telemetry process. Every technology discussed above deemed to provide significant efficiency increases, represents a significant shift in the way business is done. Some of these will not be too difficult to implement because they are hardware based, and once the hardware is built and tested people will accept and use it.

However, the one area that has the most potential for efficiency increase is also the most difficult to implement – better requirements analysis. This is due to a couple of reasons. First, automating requirements analysis is a difficult problem. This means that it will continue to require ongoing human effort. Second, as any software engineer can tell you, getting accurate, unambiguous requirements from the user is also a very difficult problem. The reference here to software is intentional. From a telemetry point of view, measurand lists and sample rates are the requirements for a piece of software. They are what are used to program the telemetry systems for a particular test. Asking for efficient use of spectrum is equivalent to asking for an efficient program to be developed every time a test is done. This is analogous to asking for efficient software requirements for a sequence of hundreds of related, but different programs to be delivered as quickly as two or three a day. Getting people to provide this is truly going to require a dramatic and very difficult change in the way people think about flight test requirements.

BIBLIOGRAPHY

[1] Advanced Range Telemetry (ARTM) website, http://tsi-elad.edwards.af.mil/artm/.

[2] Chalfant, Timothy A., Straehley, Erwin H., and Switzer, Earl R., *Advanced Ranged Telemetry (ARTM) Preparing for a New Generation of Telemetry*, Proceedings of the International Telemetering Conference (ITC), Vol. XXXII, (1996) paper number 96-24-2.

[3] Jones, Charles H., *Telemetry Mapping Algorithm Project (TMAP) Efficiency and Feasibility Study Report*, International Test & Evaluation Association Workshop, Lancaster, CA, http://www.edwards.af.mil/itea/papers, (March 1998)

[4] Jones, Charles H., and Gardner, Lee S., *Complexity of PCM Formatting*, Proceedings of the International Telemetering Conference (ITC), Vol. XXXIII, 1997 paper number 97-10-5.

[5] Telemetry Group, Range Commanders Council, IRIG Standard 106-96, Telemetry Standards, Secretariat, Range Commanders Council, U.S. Army White Sands Missile Range, NM, (May 1996).