

# Multiprogram Transport Stream: Management and Switching Solutions

By Jean Macher and George F. Anderson III



Jean Macher



George F. Anderson III

*Terrestrial DTV transport stream management can be difficult at times. Broadcasters are now required to manage network-provided multiple-program standard-definition transport streams for multicast during part of the broadcast day, locally encoded content that can be added to the bouquet of MPTS services, and high-definition content. This paper explores the technical challenges of managing the complex bouquet of programs confronting the terrestrial DTV broadcaster, with a focus on how to switch programs in the ATSC transport stream.*

Transport stream switching processes are fast becoming a necessity rather than a desire. Digital television (DTV) has progressed beyond the single-program transport stream (SPTS) of SD (standard definition) and HD<sup>1</sup> (high definition) to a sophisticated array of the multiple program transport stream (MPTS). Today's DTV transmissions require management of content, format, and a mix of SD and HD streams. Managing various programs through manually controlled switching systems lends itself to rudimentary remote control via general purpose interfaces (GPIs). Exploration of automation interfaces tailored to these requirements is still in its infancy. Program services require multiplexing combined with accurate PSI/PSIP<sup>2</sup> table management. Services selected for inclusion in the ATSC MPTS for transmission require oversight to insure that they are neither over- nor undersubscribed. Statistical multiplexing is beneficial when the stat pool contains both SD and HD content, however, it is only part of the solution.

## Transport Stream Creation in a Broadcast Environment

The implementation of DTV in the broadcast station is unique, because a single MPEG-2 transport stream with a bandwidth of approximately 19.4 Mbits/sec carries different mixes of HD and SD content during the course of a day (a typical scenario is to broadcast one HD program during primetime and multiple SD programs the rest of the day).

Today's DTV content distribution comes in two flavors:

- The Public Broadcasting Service (PBS) provides two distinct and separate satellite feeds to member stations: an MPEG-2 multiple-program transport stream consisting of four SD programs (Schedule X,

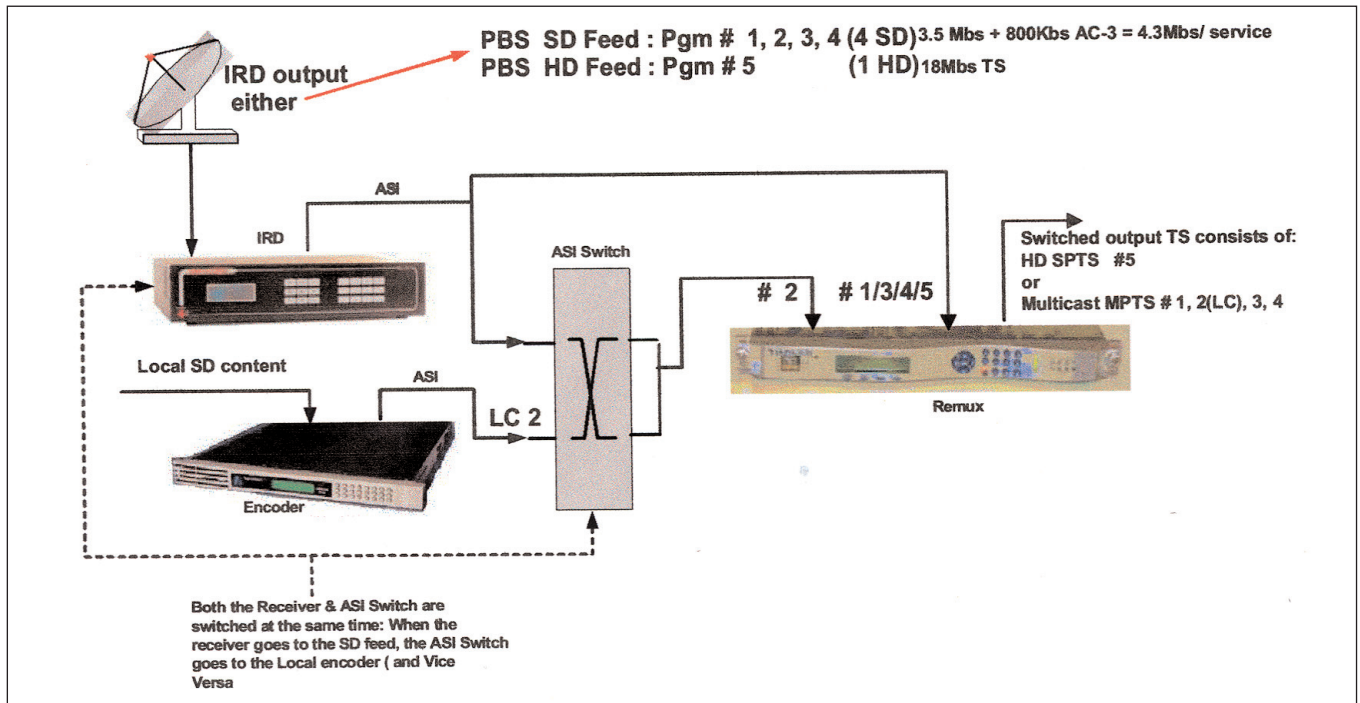


Figure 1. Simple transport stream switching.

Schedule X delayed, PBS Kids, and PBS University); an MPEG-2 single-program transport stream consisting of a single high-definition 1080i program at approximately 18 Mbits/sec.

- Commercial television networks provide a satellite delivered 45 Mbit/sec SPTS with one HD program. The source of SD material remains the legacy NTSC feeds to affiliates.

From a pure MPEG-2 standpoint, the broadcaster will use either existing pre-encoded material, keeping it in the compressed domain, or encoded uncompressed material. The pre-encoded material, if available, comes from live feeds or is stored on servers; the uncompressed material is native baseband, or results from a decoding process.

Public stations have the option of selecting from several configurations for transmission:

- Passthrough of the multicast MPTS.
- Passthrough of the HD SPTS.
- Drop SD programs and add new local programs to the multicast.
  - Decode all of the SD programs and choose from the bouquet, which programs they want to re-encode for transmission.
  - Decode the HD 1080i services, re-encode at a lower bit rate, and possibly cross-convert to one of the

other HD formats (720P) to allow for transmission of both HD and additional SD programs.

Commercial stations need to decode/re-encode their HD services feed to fit into the ATSC 19.4 Mbit/sec standard. In the rest of this document, the PBS distribution scheme will be used as an example, however, the issues and solutions described are not specific to public stations. In the context of program switching, PBS MPEG-2 feeds are equivalent to any server playing back HD or SD material, and thus apply to any kind of station. The five options listed, all require some level of PSIP management. Each station must generate its own PSIP or rebrand existing information in order to provide tuning, branding, and electronic program guide (EPG) information to the set-top box (STB).

## Switching Programs in a Transport Stream: Why and How?

### Background

There are two obvious reasons why broadcasters need to switch programs or content in an ATSC transport stream. First, HD material is only available part of the day and consumes much more bandwidth than SD material, so a broadcaster is likely to switch between programs on a regular basis. Second, local program-

ming is the key differentiator, thus, broadcasters will often replace national content (coming from a distribution feed) with local content.

### ***Program-Switching Definitions***

Switching programs is only an issue when performed in the compressed domain, that is, the MPEG-2 transport stream, without decoding back to baseband. Switching in the uncompressed domain has been available for quite some time and comes with good format conversion (e.g., SD to HD upconversion), advanced transitioning (e.g., fade-in/out), and transparent logo insertion of the analog world.

### ***Program Splicing***

One scenario in the compressed domain is when program content is switched at the elementary stream level while the program definition remains unchanged: the video/audio PIDs and the overall PSI/PSIP description of the program is the same; only the video and audio contents are switched. This stream manipulation is known as splicing, or sometimes digital program insertion (DPI).

The cable industry presently employs DPI technology to insert advertisements into the transport stream. These available slots are part of the cable carriage agreement between a cable company and a broadcaster. The Society of Cable and Telecommunication Engineers (SCTE) standards 30 and 35<sup>3</sup> provide both the foundation and the protocols for implementation of splicing into the DTV transport stream. To fully enjoy these functionalities of splicing within the transport stream, content providers need to adopt these protocols. This implies that encoder, multiplexer, and video server manufacturers also need to implement splicing protocols.

It is the authors' belief that the next generation of products and services will have these DPI splicing protocols implemented.

### ***Program Add/Drop***

Another switching scenario in the compressed domain is when a program is replaced with one containing a different definition in the transport stream—the video/audio PIDs are different, as well as program names, numbers, and all the PSI/PSIP parameters. In fact, the program can even be replaced by more than one new program.

Switching here is at the transport packet level and is

similar to the add/drop of a remultiplexer, where one program is removed from the transmission while one or several new ones appear in the broadcast transport stream (there is no concern of video/audio transition as in splicing). The key issue, however, is that in a broadcast environment, a program disappears only for a limited time—HD goes off after primetime and is replaced by multiple SDs, but returns during the next primetime. In fact, some programs in the ATSC transport stream keep switching back and forth between an active and inactive state. Even when inactive and not physically present in the transport stream, the ATSC program keeps a PSIP description at all times and is thus described in the EPG displayed by the STB.

### ***Statistical Multiplexing of all Active and Inactive Programs***

Another approach to program switching is to return to baseband for all the program content available and completely encode or re-encode the selected bouquet of programs. A broadcaster has to make this choice if pre-encoded material is not available, if it is not suited for ATSC transmission (e.g., the HD bit rate is too high), or if he wants to control and potentially modify any single video/audio encoding parameters and fine-tune his ATSC transport stream. Investing in an encoding system is costly for the station; it brings back the advanced video processing features available in the uncompressed domain prior to encoding.

Also, statistical multiplexing, if used, can provide an interesting alternative to the active/inactive program issue discussed above. Again, consider the typical example of four SDs switching back and forth with one HD. The five programs can belong to a single stat pool of approximately 19.4 Mbit/sec ATSC bandwidth. The game is then to feed the encoder with a still picture (or slide) while a program is inactive. The bit rate of an inactive program thus drops naturally to a minimum, leaving the bulk of the bandwidth available to the active programs. An advantage of this model is that problems concerning configuration switching no longer exist, because all five programs are physically present in the ATSC transport stream at all times. The drawback is that transmitting a slide in order to explain that a program is inactive is redundant with an existing PSIP mechanism and consumes bandwidth that could be used for other content.

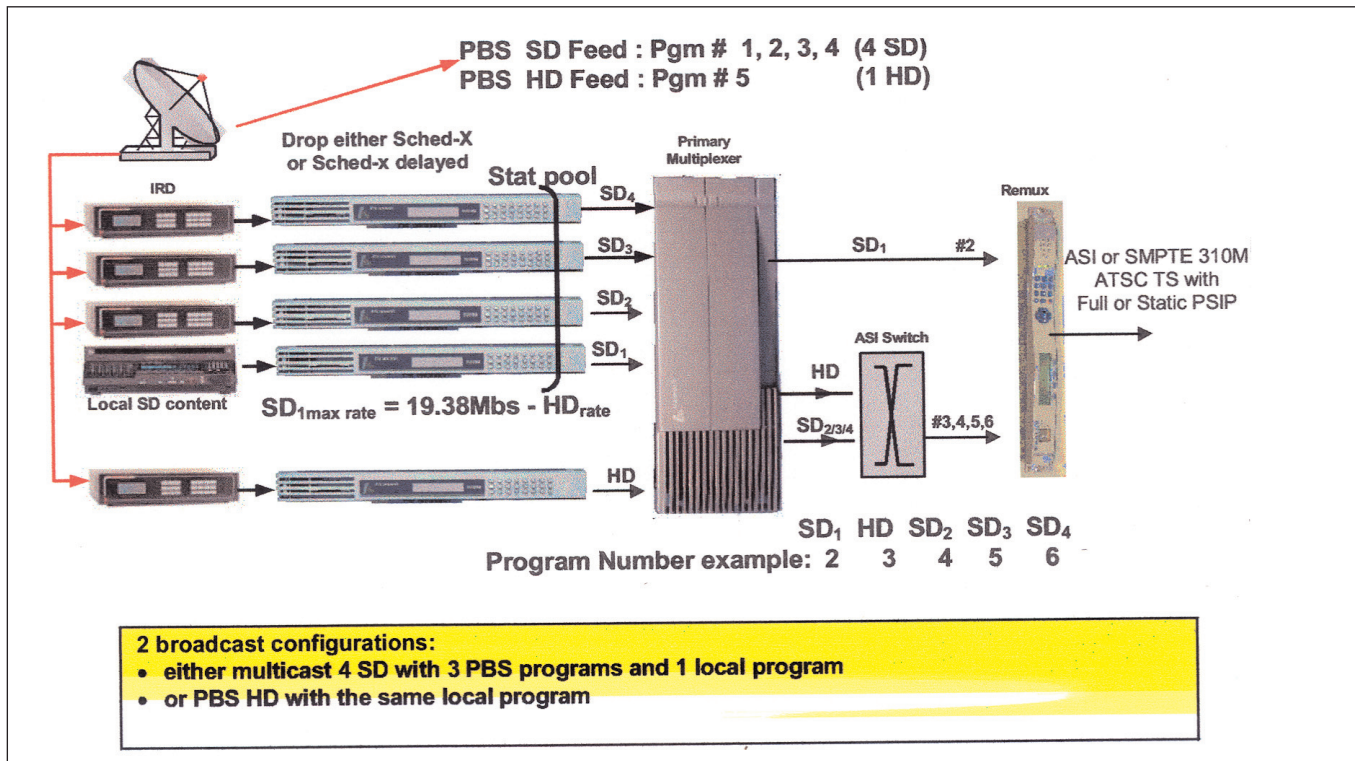


Figure 2. Decode/re-encode with stat mux and ASI switching.

### Real-Life Program Switching Solutions

The following examples are real implementations of ATSC transport stream management. They are different in cost and features, and mix the different models described above.

### Passthrough with Simple Transport Stream Switching

Fig. 1 shows a PBS station with HD passthrough and added local content to the SD multicast. Older integrated receiver decoders (IRDs) provided by PBS have a single receiver, tuned to either the SD multicast MPTS or the HD SPTS. The asynchronous serial interface (ASI) output from the IRD feeds one input of the ASI switch and one remux. The remux is configured to pass program number 2 on its first input and numbers 1, 3, 4, and 5 on its second input, thus dropping either schedule X or its delayed service (depending on the local time zone). The output from the local SD encoder is configured to match the bit rate of the replaced program.

The ASI switch and the IRD are switched at the same time, either manually or by automation, if available and practical. When the IRD goes to the multicast MPTS feed, the ASI switch goes to the local encoder,

resulting in a new MPTS consisting of the three elected PBS programs minus the discarded service that is replaced by the local SD content encoder. When the IRD goes to the HD feed, the ASI switch goes to the IRD, and the remux receives the HD program on both inputs. The remux automatically passes only the HD program in the final ATSC transport stream and injects static PSIP or receives comprehensive PSIP information (including the EPG) from a PSIP generator. This low-cost solution takes advantage of existing PBS feeds and doesn't require a full-blown encoding system. It is flexible enough to provide local content insertion and can easily integrate with an existing master control.

### Customized Decode/Re-encode Architecture

In Fig. 2, full decoding of the entire multicast and HD PBS feeds back to baseband (259M and 292M, respectfully) gives the broadcaster control over the entire bouquet of programs, including, but not limited to program look and feel, store and forward requirements, encoding formats and bit rates, statistical multiplexing, and logo insertion. The primary multiplexer is fed with all the encoders and controlled for statistical

multiplexing. The first choice the broadcaster has to make is how low the HD program can be encoded; then he can deduce the number of SD programs that can be added to the HD to fit in the ATSC 19.4 Mbit/sec bandwidth.

In this paper, the HD is encoded at 14 Mb/s, which leaves room for one more SD encoded at about 5 Mb/s. The HD program is not part of the statistical multiplexing process; only the four SD programs, which share a stat pool of about 19 Mb/s. The broadcaster decides to have the main program SD<sub>1</sub> on the air at all times, sharing primetime with the HD program and non-primetime with the three other SD programs. The primary multiplexer generates three MPEG-2 transport streams (each with the relevant PAT/PMT tables) on three independent ASI outputs. One output carries SD<sub>1</sub> and feeds the first remux input; one carries the HD program and feeds the first input of an ASI switch; and another, with SD<sub>2</sub>/SD<sub>3</sub>/SD<sub>4</sub>, feeds the second input of the ASI switch. Finally, the second input is fed with the output of the ASI switch.

The remux is configured to get SD<sub>1</sub> from input 1 and HD/SD<sub>2</sub>/SD<sub>3</sub>/SD<sub>4</sub> from input 2. Because the ASI switch passes either the HD program or SD<sub>2</sub>/SD<sub>3</sub>/SD<sub>4</sub>, by controlling that switch the ATSC transport stream carries either HD/SD<sub>1</sub> or SD<sub>1</sub>/SD<sub>2</sub>/SD<sub>3</sub>/SD<sub>4</sub>. In addition, the independent SD<sub>1</sub> feed between the primary multiplexer and the remux guarantees that SD<sub>1</sub> suffers no disruption during the transition from one configuration to another. The system described in Fig. 2 seems a viable tradeoff between statistical multiplexing efficiency and HD-to-multicast switching flexibility.

**Passthrough with PSA Insertion**

In this Fig. 3, the station does a simple passthrough of the PBS feeds, rebrands the incoming PSIP, but also attempts to insert some short local identifiers or

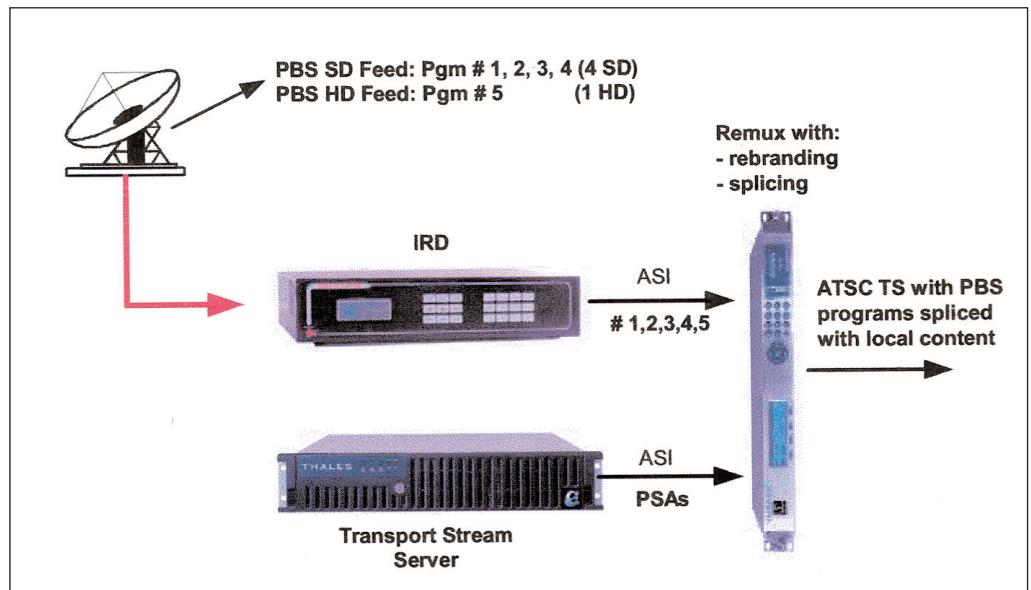


Figure 3. Passthrough with program content splicing.

public service announcements (PSAs) in place of the PBS programs. From a viewer standpoint, there is need to tune to a different program. The PSA is just a sequence in the middle of the PBS program's video and audio content. The switching is done at the content level; this is a splicing scenario. The first input of the remux receives the PBS feed through the IRD. The second is a transport stream server with pre-recorded PSA sequences (small PSAs could be directly loaded on the remux).

The system described in Fig. 3 represents a cost-effective alternative to logo insertion (whether in the compressed or uncompressed domain) for a station that doesn't have an encoding system and relies on network feeds passthrough.

**Program Switching and PSIP Management**

**Inactive Programs Issue**

Program content splicing is transparent for the PSI/PSIP tables. The description of the program conveyed in the PAT/PMT/VCT tables remains the same before and after the splice. Switching programs in the ATSC transport stream, however, has impact on the PSIP management beyond the usual creation and rebranding considerations.<sup>4</sup>

As mentioned earlier, during primetime an HD and an SD program are broadcast, and four SD programs

are broadcast the rest of the day. This introduces the notion of active and inactive programs during the course of a broadcast day. PSIP signals an inactive program via the hidden flag of the virtual channel table (VCT).

Today's PSIP generators can schedule this flag for each program in the ATSC transport stream, however,

because the encoding system must also be reconfigured (via a simple ASI switch in this scenario), there is a risk that the VCT will not be consistent with the content of the final stream. Typically, this could prevent viewers from tuning to an active channel that is declared inactive.

The solution consists of handling the hidden flag in

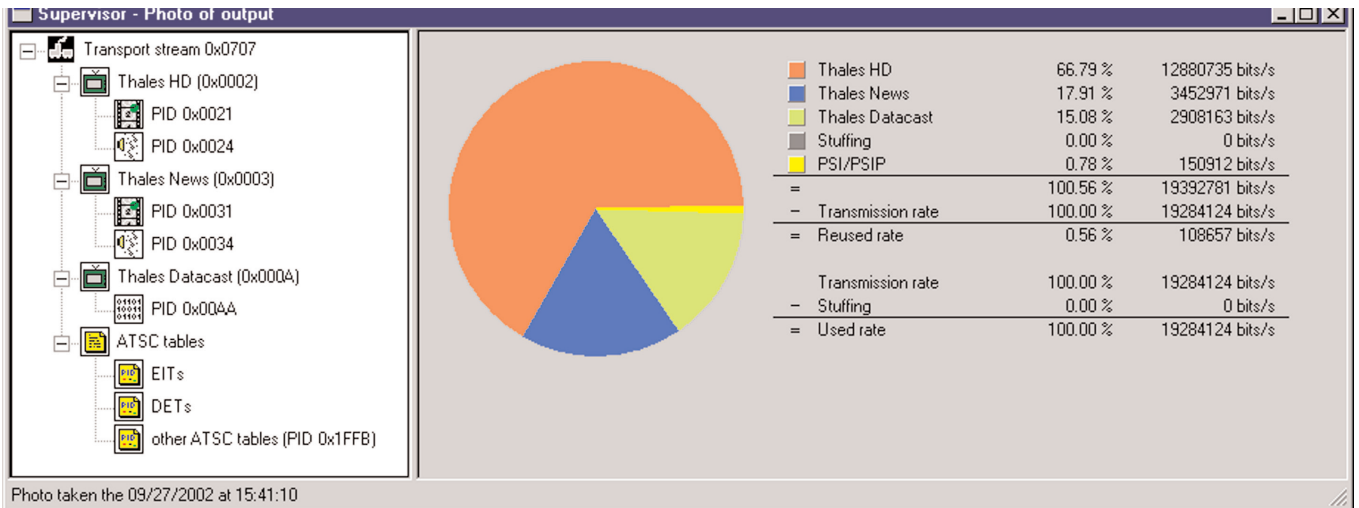


Figure 4. MPTS of 1 x HD + 1 x SD + 1 x datacast.

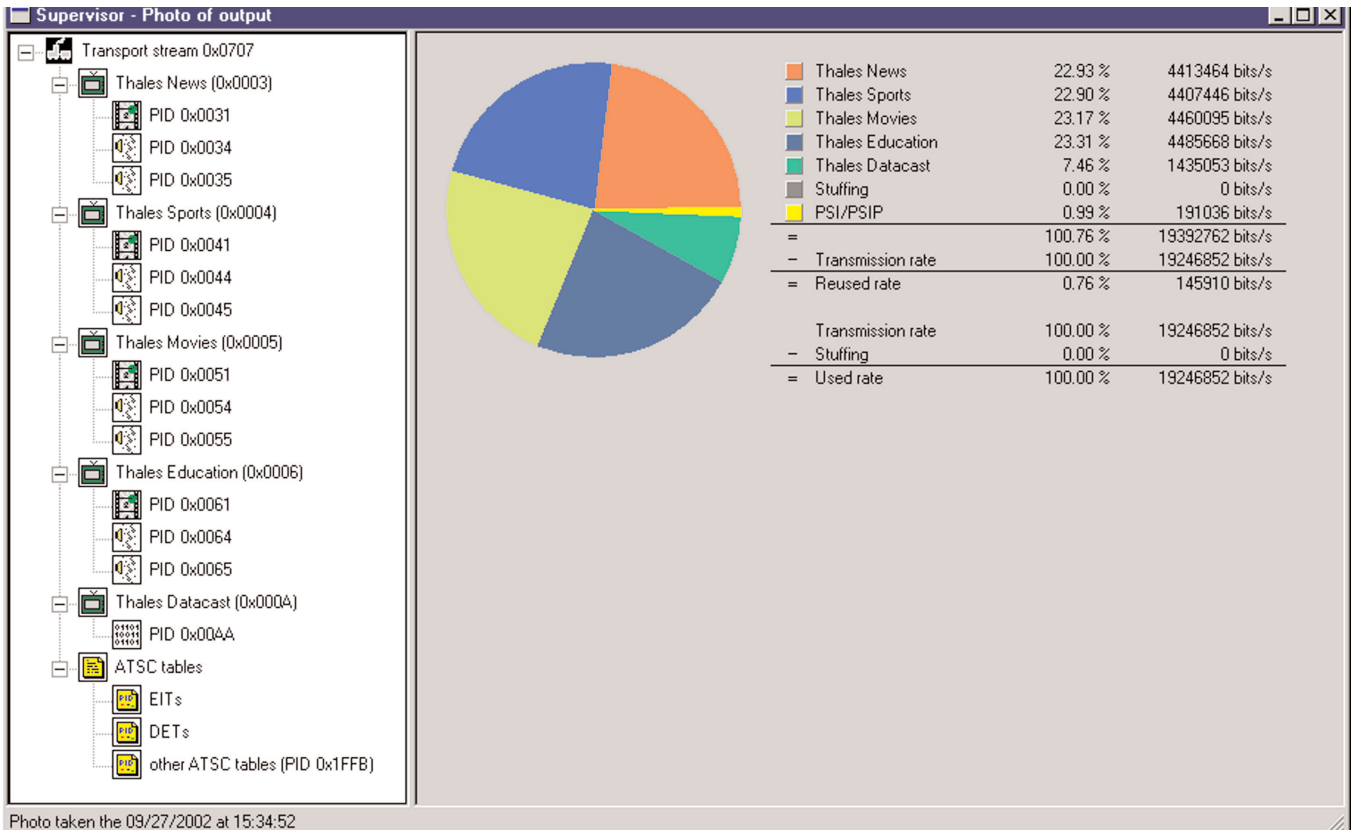


Figure 5. MPTS of 4 x SD + 1 x datacast.

the equipment that builds the final ATSC transport stream (i.e., the remux in this case) because it knows exactly which programs are currently active. The remux can analyze the incoming program association tables (PAT) to determine which programs are active, and then, for each program declared in the VCT, update the hidden flag accordingly.

### Implementation Example

In Fig. 2, the remux runs a single configuration by looking for program number 2 on its first input and numbers 3, 4, 5, and 6 on its second. All these programs are listed in a VCT that is either generated by the remux (static PSIP scenario) or comes from a PSIP generator. Programs come and go as a result of upstream switching, and the remux updates the VCT and generates a consistent ATSC transport stream.

Figures 4 and 5 show a snapshot (program tree view and bit rate pie chart) of the ATSC transport stream coming out of the remux at two different times during a broadcast day. Again, no configuration change, other than at a simple ASI switch level, is required to obtain this dynamic behavior.

### Conclusion

Program switching in the ATSC transport stream is a unique challenge for U.S. broadcasters that switch back and forth between SD and HD material to fit the 19.4 Mbit/sec bandwidth. Public broadcasters benefit from two PBS MPEG-2 transport stream feeds, and can elect to return to baseband or stay in the compressed domain to tackle their program-switching needs. Switching schemes taking advantage of simple ASI routing are reasonable solutions for today's complex switching requirements. Applications of these solutions are available to both small and large broadcast stations.

### Notes

1. For the purpose of this paper, HD format refers to one of the following: 480P, 720P, or 1080I.
2. "Program Specific Information," as defined in ISO/IEC 13818-1 and "Program and System Information Protocol," as defined in ATSC A/65.
3. Splicing and PDI insertion as described in SCTE 30 & 35.
4. PSIP creation and rebranding are out of the scope of this paper, but please refer to the "Real World PSIP Solutions" paper presented by Thales Broadcast & Multimedia during NAB 2002, for a comprehensive overview on the subject.

## THE AUTHORS

**Jean Macher** joined Thales Broadcast & Multimedia in 1998 and relocated from France in 2000, taking on the role of digital television product line manager for the U.S. In France, he was in charge of the company's service information (SI) processing product line, and supervised the specification and adaptation of remultiplexing and SI processing products to U.S. markets. Macher previously worked as a project engineer for the company's U.K. Digital Terrestrial TV Project (remultiplexing and SI management for the BBC, CTI, On Digital), developing technical proposals and SI architecture specifications. He holds a Diploma of Engineering from the Superior Institute of Electrical Engineering of Paris (ISEP), a major French computer and electronic school (equivalent to an MSEE).

**George F. Anderson III**, is sales manager of digital services for Thales Broadcast & Multimedia. He started at Thales in January 2000 as sales manager for Comark Digital Systems (CDS), where his responsibilities included the promotion and sales of all the digital services necessary to interface the station's existing NTSC master control to high definition, SD and HD encoding, and datacasting capability, as well as to the station transmitter link distribution system.

Anderson has many years of experience in broadcast sales and media management, including the movement of media both inter- and intra-facility, and providing media management solutions for the emerging digital media market, previously known as broadcast television. He holds a degree in electrical engineering with a minor in marketing from the University of Alabama. Anderson is a past Chair of the Philadelphia Section of SMPTE, and also served as Eastern Region Governor in 1999.