THE BROADCAST BRIDGE
Connecting IT to Broadcast

Using Configurable FPGA's To Deliver Integrated Flexibility, Scalability, And Resilience

Essential Guide

EG
ESSENTIAL GUIDES

Introduction

By Tony Orme, Editor at The Broadcast Bridge

Advances in hardware design are delivering an improved immersive and augmented viewing experience for audiences as software and hardware combine to provide the ultimate in flexibility and scalability.

Without doubt, software has had a major influence on the broadcast industry in recent years. Although software has existed in application specific products for as long as we can remember, it's only since we started looking seriously at COTS has it had such an impact on not only workflows, but video and audio processing too.

As pixel counts and frame rates increase, then so does the complexity of our new infrastructures. Pure on-prem and cloud systems are not always feasible leading broadcasters to again look at hardware as a solution. It's simplicity of operation and predictably low latency often shines through. However, the seemingly static nature of hardware has until recently made broadcasters question the wisdom of procuring it. But a new breed of flexible hardware is delivering the best aspects of hardware and software to deliver predictably low latency and flexibility.

FPGAs have been used in broadcasting for many years but only deep within a vendor's product. A production switcher may have tens of FPGAs in its design, but their functionality has not been directly available to users. They are particularly attractive as they represent a massive array of hardware logic that can be configured in software to provide truly parallel hardware solutions.

When combined with software management systems, directly configurable FPGA resource provides massive on-prem flexibility, scalability, and resilience. But this is not the only benefit as a further advantage becomes apparent when we think in terms of connectivity.

One of the limitations of pure-cloud and on-prem infrastructures is that the only physical connectivity available, especially with public cloud, is IP over ethernet, or another IT specific transport stream for IP.

FPGA flexible hardware systems abstract the underlying transport stream to improve the number of options a broadcaster has of distributing their media. Few broadcasters can build a complete greenfield IP site and so will migrate to IP one workflow at a time. In doing so, they create a hybrid system that has the potential to be incredibly complex. Even when a broadcaster may think they will always keep their SDI infrastructure, as IP solutions develop, there may well be compelling reasons to adopt some IP workflows.

As FPGAs are effectively huge programmable parallel processing devices, streaming media between SDI, AES, or IP is a task well within the constraints of the device to further improve system flexibility while keeping latency predictable. Combined with a management software control system, assigning resource between varying transport streams becomes a drag-and-drop operation.

Viewers are making unprecedented demands on their media entertainment suppliers meaning flexible systems are key to maintaining modern broadcast infrastructures. Also, as the systems become more advanced, vendors are rising to the challenge and abstracting the underlying complex technology so engineers and technologists can get on with delivering the workflows that best meet the needs of their viewers.

Tony Orme.

In the same way automobile manufacturers innovated their vehicles to provide fuel injection and software management to improve efficiency and reduce complexity, broadcast vendors are now delivering truly flexible, scalable, and resilient hardware-software infrastructures. Not only do these systems keep latency predictable, but they also abstract away the complexity of the underlying transport stream allowing broadcasters to bring workflows together and focus on delivering for their viewers

Tony Orme
Editor, The Broadcast Bridge

Using Configurable FPGA's To Deliver Integrated Flexibility, Scalability, And Resilience

By Tony Orme, Editor at The Broadcast Bridge

Although IP and cloud computing may be the new buzz words of the industry, hardware solutions are still very much in demand, especially when taking into consideration software defined architectures. And in addition, a whole new plethora of FPGA based configurable systems are providing flexibility, resilience, and scalability.

Broadcast facilities have traditionally been designed using standalone devices that perform a particular task such as a standards converter, proc-amp, or production switcher. However, the new breed of software defined hardware systems allow generic hardware to be assigned to different tasks depending on the software loaded into it.

This leads to transport stream agnostic workflows that can work with SDI, AES, ST2110, and many other protocols, all with simplicity and ease.

This integration of transport stream agnostic workflows is empowering broadcasters to decide which transport technology they want to use, progress to, or stick with. For example, a broadcaster may want to keep SDI/AES but have the option of working with IP or Cloud services later.

With FPGA systems broadcasters get the best of all worlds, that is, flexible and simplified workflows with very low latency. Although FPGAs may be programmed with software files, they are intrinsically hardware devices that work to very tight time constraints and keep latency to sub nanosecond levels across the whole processing fabric. In essence, they deliver the ultimate parallel processing solution.

Advancing Technology

Broadcast engineers and technologists have had a lot to deal with over the past twenty plus years. Many have had to transition from analog NTSC and PAL, to SDI, widescreen, HD, then HDR and even UHD and 8K. And that's before we start talking about compression and IP. Transitioning directly to IP may well be the end game in terms of the technological evolution of broadcasting, but for many this change must be managed carefully

Television has traditionally been risk averse, and in part this is due to the eye watering amounts of money that are involved with live sports productions and premium movie rights. But it's also due to the ongoing technological development of video and audio processing. Broadcasting is like an innovation sponge, no matter how fast scientists and researchers deliver new high-speed circuits, broadcasting always demands more. It was not so long-ago that HD was state of the art with its 1.5Gbits/s circuits, now we're talking in terms of 12Gbits/s SDI for 8K and beyond.

IP delivers some outstanding benefits, but not all broadcasters need these. And even those that do, may find that not every workflow lends itself well to a full IP migration. There are many processes that still benefit greatly from a hardware approach especially when we start to consider latency. One of the challenges we have with IP is that software processing is inherently indeterminate in terms of latency, but this may or may not be an issue for the broadcaster. The good news is, that hybrid hardware systems can deliver the best of both worlds.

Latency Challenge

Hardware has always been dominant in the broadcast industry as it was generally the only method we had of processing real time video and audio with very low and predictable latency. But as computer processing, storage and network distribution have all increased in speed and reliability, it is now possible to process video and audio in real-time. The challenges we have are that software-based systems are both complex and latencies are indeterminate.

Looking at this it would appear we either opt for inflexible and static hardware systems that have determinate latency, or highly flexible and scalable software systems that suffer from highly variable latency. However, the ongoing technology innovation is reducing the latency in software and improving the flexibility and scalability in hardware. A combination of the two could prove the perfect hybrid solution.

Another major trait of broadcaster engineers and technologists is they like to keep their options open and not back themselves into a technology corner. This is achievable when we look at the new breed of FPGA hardware solutions. Not only do they provide all the benefits of hardware processing, but they also abstract way the transport stream so that the broadcasters are not tied to IP, SDI or AES etc. As their needs change and the technology develops, broadcasters can move more to IP or keep with SDI and AES. Fundamentally they significantly increase their options.

FPGA Solutions

With modern FPGA approaches to system design broadcasters now have the best of both worlds, they have the flexibility and scalability of software with the low and predictable latency of hardware.

FPGAs have advanced massively in recent years. Originally, they were seen as a flexible method of reducing hardware real-estate on circuit boards, but as their features and capacity increased, then so has their usefulness.

If we look at a simple video digital filter, such as a Finite Input Response (FIR) filter, then the main components are delay lines, adders, and multipliers. Using discrete components would take an enormous amount of real-estate on a circuit board even using surface mount components. Impedance matching the tracks between the components would be a real challenge at HD and UHD SDI speeds. Maintaining high manufacturing yield rates at these frequencies is heavily impacted by the number of components needed on the circuit board. However, high-capacity, and high-speed FPGAs can facilitate FIRs relatively easily, even at 4K and UHD video rates.

We could use DSPs instead of FPGAs as they have the flexibility of software, but they do suffer from limited resource, especially when we consider video rates. Furthermore, many more devices with external memory would be needed to replicate the equivalent high-end FPGAs.

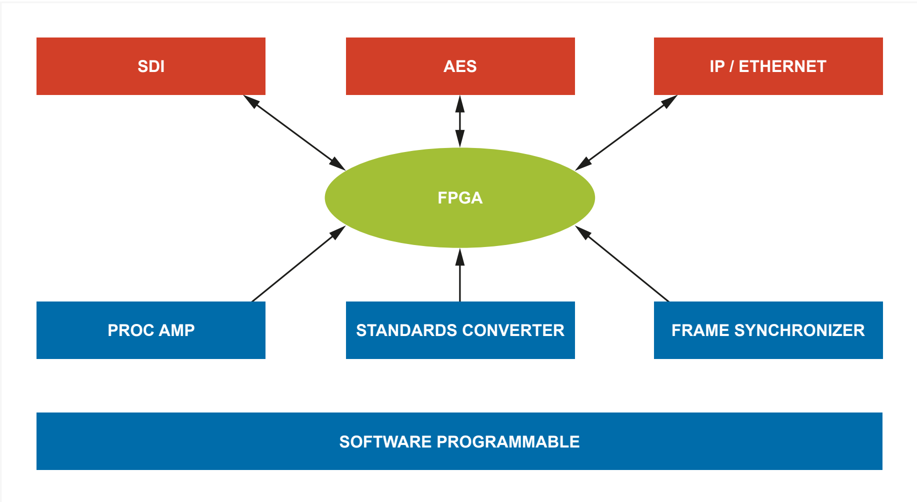

Figure 1 – FPGAs can provide high speed inputs and outputs for SDI, AES, IP Ethernet and Fiber, and their functionality can be programmed using software to deliver all the advantages of high-speed and low-latency hardware, along with the benefits of scalability and flexibility found with software.

Flexible Hardware Resource

FPGAs are a programmable device that use binary files to interconnect resources within the device. Fundamentally, these consist of gates such as AND and XOR, and from these two gates any of the Boolean functions can be replicated, which forms the basis of any digital design.

Another alternative for processing high-speed video is to use ASICs (Application-specific integrated circuit). These maintain the speeds of FPGAs (and often faster) but are fixed in their functionality. That is, if an ASIC is designed to provide MPEG2 compression functionality then that's all it can ever do. Forward thinking ASIC designers will build into the silicon registers that can be updated from an external CPU so some of the compression parameters can be changed, but other than this, the design is locked-in during manufacture. The big advantage of ASICs is that long manufacture runs soon reduce the cost making consumer equipment affordable.

FPGAs are significantly more expensive than ASICs or the equivalent DSP solutions, but they are orders of magnitude more flexible than any other solution on the market. Not only are FPGAs programmable, but their configuration can also be changed on-the-fly. This allows an array of FPGAs on a circuit board with some external memory and interface ICs, to act as a massive programmable hardware resource.

Not only do FPGAs provide Boolean logic resource, but they also have built into their fabric features such as block memory, DSPs, CPUs, and multipliers. In effect, an entire hardware ecosystem exists on a single chip. Also, high-speed interfaces are provided so that multiple devices can be interlinked to further increase their capacity.

Although many features such as memory, DSPs, CPUs, and multipliers can all be replicated in Boolean XOR-AND logic, doing so is inefficient. For example, a simple D-type single bit memory function can be provided using five NAND gates, which in turn will require three XOR gates and one AND gate for each NAND gate, resulting in thirty XOR gates and ten AND gates for a simple 10-bit luminance single clock delay line. This is used in a single tap in a video FIR filter, which in turn could consist of twenty or thirty taps. And this is before we start looking into the complexity of adder circuits, multiplexers, and multipliers.

Software Hardware

These hardware resources take the benefits of ASICs but maintain the flexibility of software programming. FPGAs are available in all different shapes and sizes and are also defined by the amount of hardware resource they have available to the design engineer.

Programming an FPGA is often a three-stage process: modelling, simulation and verification, synthesis and placement. All these stages are performed in software and the output is often a formatted binary file that is loaded into the FPGA during boot or at any other time under the control of the system.

Modelling is the operation of designing the process the engineer is building and is often facilitated using languages such as Verliog and VHDL. Both these are considered High Description Languages and in appearance are not dissimilar from procedural programming languages. Simulation and verification provide offline testing of the design where data samples are presented, and the outputs are verified against the expected output. For example, an FIR is highly determinate, and we know what the output values should be given a known input. The final stage is synthesis and routing where the binary file is built and programmed into the FPGA.

All these processes can take an incredible length of time, especially when the designs become complex and are therefore often divided into smaller test benches. But once the design of the function is complete, the final software file is downloaded into the FPGAs within a matter of milliseconds.

What is really compelling about FPGAs is that once the circuit board hardware is designed and built, the rest of the implementation is based on software. This provides untold flexibility for vendors as they can literally make the hardware do anything they want (within the limits of the resource). And this is a real opportunity for broadcasters.

Proven FPGA Technology

Although FPGAs have been used for many years inside broadcast hardware designs, and are therefore proven technologies, it is only recently that arrays of FPGA ICs have become available as stand-alone ecosystems to facilitate dynamic and scalable resource for broadcasters. For example, a single card could be programmed to be a proc-amp, but on the next day be reprogrammed by downloading a binary file, to be a standards converter. This flexibility is something we've never seen before in broadcasting when considering the very low latencies involved. It's possible for software COTS to deliver this, but the latencies are variable and unpredictable, and the systems are incredibly complex, which is often an issue for live productions. FPGA arrays are low and latency determinate, and relatively easy to operate.

FPGA IC arrays provided on a single circuit board can then be replicated many times in a rack frame. With the appropriate management software, the functionality is effectively abstracted away from the underlying hardware, including the transport stream. This delivers incredible opportunities for flexible and scalable operation for broadcasters, especially when considering the potential for the multitude of licensing models. For example, using the pay-as-you-go model, centralized licensing repositories could be linked into the vendors management software to make available modular functionality, such as proc-amps, embedders, and frame-synchronizers, to name but a few.

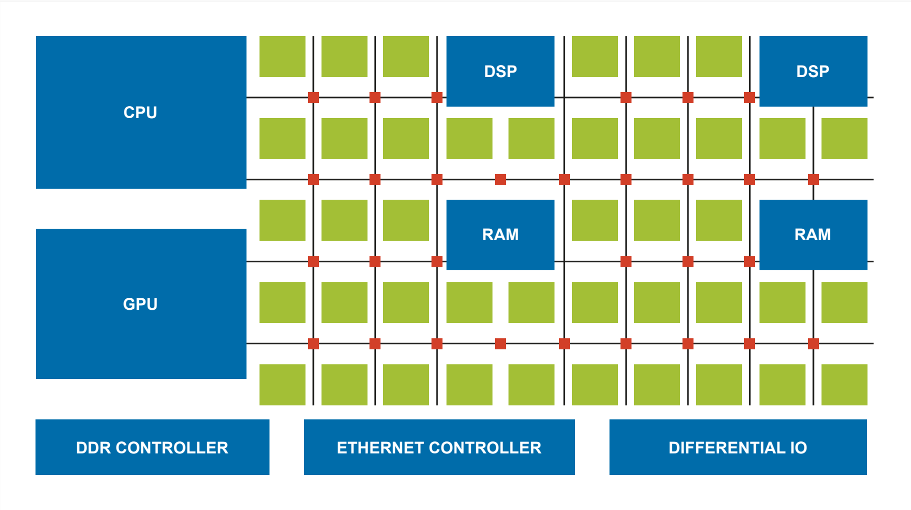

Figure 2 - An FPGA consists of tens-of-thousands of hardware gates and functions that can be programmed allowing many different operations to be provided such as proc-amps, standards converters, and frame-synchronizers. Also, many FPGAs can be connected with high-speed differential pair busses to facilitate low latency signal processing across multiple FPGAs.

The really exciting aspect of this initiative is that when the rack of FPGA cards has been procured, all the operational functionality is provided by the vendor using software files. These can be updated and managed by the vendor so the broadcaster can focus on building their specific studio solution without having to worry about software versioning or configuration. Furthermore, by increasing the number of racks, the available resource scales appropriately. Therefore, when an engineer is designing or expanding their facility, they can spread their estimated functionality requirements over many racks knowing the detail of operation can be loaded into the FPGAs as required, thus making the system highly flexible and scalable. Futureproofing is provided without having to plan ten years ahead as more FPGA racks can be added as required.

Transport Stream Independent

This design philosophy also has some very interesting implications for the transport stream interfacing as it is taken care of by the FPGA circuit boards. The SDI, AES, ST2110 IP, and ST2022 IP protocols, and many others, are available as VHDL code libraries and so manifest themselves as physical interfaces on the FPGA. Consequently, transferring video and audio data to and from them is a relatively straight forward process as it's all taken care off in the FPGA itself. We don't need to be concerned with interface equipment to convert between the various transport streams, it all takes place inside the FPGA.

It's fair to say that the hardware still needs physical interfaces and connectors, but again these can be provided as an array of assignable flexible resource instead of being statically dedicated to specific tasks, thus further improving flexibility and scalability.

Another interesting aspect of this design philosophy is that much of the video, audio and metadata signal routing takes place within the confines of the rack of FPGA circuit boards through high-speed back planes, not only does this keep latency low, but another positive side effect is that cabling is significantly reduced.

Although cabling forms the core of any infrastructure, it has two undesirable attributes, it's heavy, and is susceptible to damage. Weight is particularly important for OB trucks and anything we can do to reduce it is a major bonus. Even where fiber is used to distribute IP, there are clear advantages to keeping the amount of fiber to a minimum, that is, the less we have, the less there is to go wrong.

Keeping the signal processing within the relatively close proximity of a rack will help maintain resilience and equally importantly, low and predictable latency. This also helps reduce the number of inputs and outputs on the central routing matrix, further keeping weight low and power consumption down.

Multiple racks with dual power supplies and diverse power routings delivers high resilience, especially when combined with a software configuration and management system that can assign the FPGA functionality on-the-fly, thus delivering outstanding flexibility, resilience, and scalability.

Conclusion

Broadcasters looking to upgrade, improve, or expand their facilities are currently presented with some very difficult decisions. In part, this is due to the influence of IP and cloud computing. However, much of the functionality broadcasters currently need are difficult and sometimes impossible to implement in IP and cloud, and this is just a natural consequence of the state of IP and cloud development at this moment in.

The good news is that the new breed of assignable FPGA arrays not only makes the delivery of flexible and scalable functionality a reality giving an outstanding compromise between hardware and software, but also abstracts the SDI, AES and IP transport streams away from the user operation allowing broadcasters to mix and match technologies with ease, and build the most flexible, resilient, and scalable broadcast infrastructure possible.

The Sponsors Perspective

Combining Hardware And Software Creates A Graphical Experience To Remember

By Todd Riggs - Director, Product Management - Hyperconverged Solutions at Ross Video

Providing an immersive video experience for fans while not detracting from the on-field action on game days is the ultimate goal.

New types of production facilities dedicated to sports venues, including control rooms located inside the stadium or ballpark, have traditionally been designed using standalone devices that perform a specific task such as a production switcher, router, multiviewer or replay system. However when you add the ability to display graphic content on a multitude of screens throughout a stadium using software-defined systems, it can be a real game changer.

A New Breed Of Software-Defined Infrastructure

Indeed, this new breed of software-defined hardware systems allow one or two operators to manage an entire system of LED screens and ribbon displays in a sports venue with far-reaching possibilities—for both fans and sponsors. The inherent flexibility is particularly valuable for multi-purpose venues that have to create, change and distribute content on the fly for a variety of clients with different needs.

Take for example Truist Park, home to the Atlanta Braves of Major League Baseball (MLB). Inside the ballpark is an HD videoboard in center field and over 32 LED boards throughout the stadium and surrounding entertainment district that require refreshed graphics for each game/event held there.

This is all managed by the "BravesVision" team that leverages a software-with-real-time hardware-processing venue control system to seamlessly trigger every display and device while simultaneously integrating closed captioning and emergency messaging. The system also relies on a 24/7/365 scheduling application operated from a production control room in the ballpark to drive every aspect of the venue experience. This includes LED displays, audio, lighting, and mechanical features throughout the district.

The solution chosen was a software-only graphics design and control system we designed called XPression Tessera that drives the multitude of displays at the ballpark and surrounding entertainment spaces. The system is operated by the BravesVision team from the hardware-based production control room.

In fact, with the tight integration between the real-time motion graphics system and a full complement of production control room technology (also from us), Truist Park became the first venue in Major League Baseball to deploy comprehensive Unified Venue Control. Content for all 32 LED screens is rendered in real-time using eight channels and operated from a centralized control room.

Supporting Multi-Purpose Venues

Likewise, California's SoFi Stadium, home to the Los Angeles Rams and Chargers NFL teams, is an end-to-end 4K HDR broadcast facility that's often referred to as the most technologically advanced sports and entertainment venue on the planet. It includes both a massive center-hung infinity screen and a translucent roof containing a network of LED lights that project images (visible by flights into and out of the nearby LAX airport).

The stunning venue features an asymmetrical roofline containing a network of IP-connected LED panels. The centerpiece is a 360-degree ovular double-sided 4K HDR screen measuring 360ft (110m) long and 150ft (46m) high. It spans 120 yards and contains nearly 80 million pixels. As a multi-purpose venue that plays host to two separate NFL teams, SoFi Stadium's production team needs tools that are flexible enough to transition from hosting an L.A. Rams game one week to an L.A. Chargers the next week and a mega concert event in-between. They now have this capability, which can easily be deployed at a moment's notice.

Skarpi Hedinsson, SoFi Stadium's Chief Technology Officer, was tasked with displaying pristine video content and 3D animations on all of the oversized LED displays, sometimes simultaneously, and sometimes on only a select number of screens. This is enabled by the flexibility of software, in combination with real-time hardware processing.

As a multi-display, real-time graphics designer and controller, XPression Tessera can drive any LED display size and resolution with pre-rendered and real-time rendered 3D graphics. A single animation can seamlessly occupy the entire 120-yard infinity board.

This flexibility means that the same graphic or animation can also be used to treat the eight primary sections of the screen at SoFi Stadium as distinct canvases, allowing operators to present a variety of looks for their content during a live event. The graphics are produced and executed from an on-site control room that contains Ross Video production technology, such as the Acuity production switcher and Ultrix routing system.

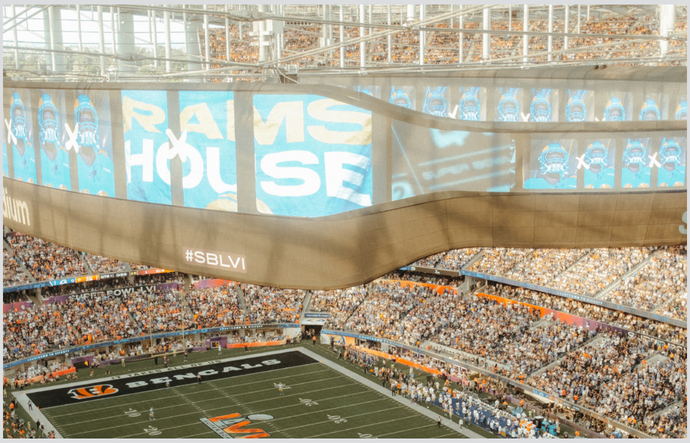

From the management of multiple internal LED displays, concourse feeds, and external displays to the control of third-party devices like video processors, LED lighting, and audio systems, the Ross Unified Venue Control Solution gives the SoFi team complete control of their production and enables them to deliver a truly immersive live experience for fans in attendance.

Using multiple SDI-based Acuity production switchers in the in-stadium control room allows the production team to cut pre-game, in-game, and post-game feeds simultaneously. When combining the switcher functionality with the software-defined Ultrix routing system, the team has an extremely powerful hyperconverged production solution that merges switching, routing, multi-viewers, and signal processing into a single platform.

In addition, a PIERO Sports Analysis tool delivers virtual first down lines, red zone markers, sponsor logos, and more to keep fans informed. A social media management tool called Inception is also used by the stadium crew to publish social media content to the ribbon boards in the stadium in real-time.

Unified Control Solution

In the UK, The Principality Stadium, located in the city of Cardiff, is the national stadium of Wales and home to the Welsh Rugby Union (WRU). The previously named Millennium Stadium seats 74,500 and, in 1999 when it opened, was the first multi-purpose sports and entertainment venue in the country to feature a fully retractable roof.

Since hosting the Rugby World Cup Finals in 1999, the Principality Stadium has staged several high-profile events, including major concerts, the UEFA Champions League Final, the FA Cup, the League Cup Final, and the Speedway Grand Prix of Great Britain.

Looking to attract more major international sporting events, the time had come to upgrade their in-venue production system and late 2021 they welcomed back fans to the venue for the first time in 18 months.

Local Systems integrator PMY Group was tasked with managing the in-venue technology upgrade project and making the stadium control room capable of producing 4K UHD content. PMY turned to the Ross Sports & Live Events group to come up with a plan.

Our Unified Control Solution was selected, after recognizing the creative, technical, and business benefits of integrating the production control room technology with the venue's LED content management system. By connecting these systems via the DashBoard venue control software, in-venue operators can trigger fully synchronized events and graphics across all in-venue displays with the press of a button.

Another benefit of the Unified Control Solution is its ability to connect with third-party data providers to keep fans in the stadium informed with the latest sports scores from around the league as well as for educational information across a variety of different sized screens.

The production control room inside The Principality Stadium features a full complement of Ross Video production technology, including a Carbonite Ultra production switcher, Ultrix router, Xpression for CG graphics, XPression Tessera to drive the LED displays, openGear cards for signal processing, and a Mira+ replay system.

Keeping In-Stadium Fans Engaged

Integrated hardware and software systems can be a powerful (and less expensive) tool to help simplify the management of a live sports production, whereby custom touchscreen control panels can be used to trigger many different macros and events. And these systems are completely customizable for each venue, supporting video formats from SD all the way to UHD. Adjusting the video output is as easy as changing an I/O card.

For sports stadium and ballparks around the world, keeping fans engaged during games with graphic displays that elevate the live experience has become a major effort to combat people watching at home. Integrating hardware and software into a single production platform to create this experience makes a lot of financial and operational sense.

For fans, a thrilling game day experience will be remembered for a long time.

THE BROADCAST BRIDGE
Connecting IT to Broadcast

Themed Content Collection

EG
ESSENTIAL GUIDES

MEDIA

ARTICLES

For hundreds more high quality original articles and Essential Guides like this please visit:

thebroadcastbridge.com

1/2023