

# Bluetooth Hi-Fi Stereo IP

Integrating the Advanced Audio Distribution Profile

▶ September 2006

Almansor Kerroum  
IPextreme, Inc.



*Integrating the Advanced Audio Distribution Profile into an existing Bluetooth product is not as straightforward as it might seem; in fact, it poses some interesting challenges.*

*In this paper, we outline those challenges and show you how to overcome them.*

## HIGHLIGHTS

- ▶ Bluetooth HW/SW Architecture
- ▶ Bluetooth Audio
- ▶ Solution for Hi-Fi Stereo
- ▶ Prototyping with Available IP

## ABSTRACT

The proliferation of portable music players has created a huge demand for wireless, hi-fi stereo quality audio devices. Such devices are now practical with Bluetooth technology thanks to the Advanced Audio Distribution Profile (A2DP). In this paper, we explore the architectural options available for implementing the A2DP in a Bluetooth-enabled device and discuss their tradeoffs. We also explain how we built an optimized solution using readily available IP, highlighting the challenges involved and how we resolved them.

## TABLE OF CONTENTS

INTRODUCTION .....	3
SPEECH QUALITY ONLY.....	3
HI-FI STEREO QUALITY .....	5
CHALLENGES .....	6
FROM SPECIFICATION TO IMPLEMENTATION .....	8
PUTTING IT ALL TOGETHER .....	11
BENEFITS.....	13
SUMMARY .....	13

## INTRODUCTION

Consumer demand for short-range wireless connectivity in devices like mobile phones, portable media players, and PDAs has led to the massive spread of Bluetooth technology. About 6 million Bluetooth ports are shipping every week and ABI Research expects that volume to rise 71% in 2006. Ease of use, interoperability, low power consumption, and wearability are the key factors to success. Further growth can be expected for medical, health, and fitness devices.

Bluetooth uses profiles to define various classes of functionality. In the audio domain, earlier versions of Bluetooth provided only telephony services, the so-called "Headset Profile," which offered only speech-quality audio. This profile was unable to satisfy the consumer craving for hi-fi stereo headsets. To support hi-fi stereo, a new profile had to be introduced, the "Advanced Audio Distribution Profile (A2DP)."

Using the A2DP (also called Bluetooth stereo), consumers now are able to stream high-quality audio from their portable devices like media players, cell phones, and PDAs to a Bluetooth stereo headset. Car buyers wish to transfer music from their portable cell phones or portable media players directly to the car stereo system. In the automotive infotainment market, Bluetooth grew by 141% between 2004 and 2005 as a study by ABI Research discovered. Those numbers are mainly based on telephony services, since the A2DP has not yet been widely deployed in car stereo systems.

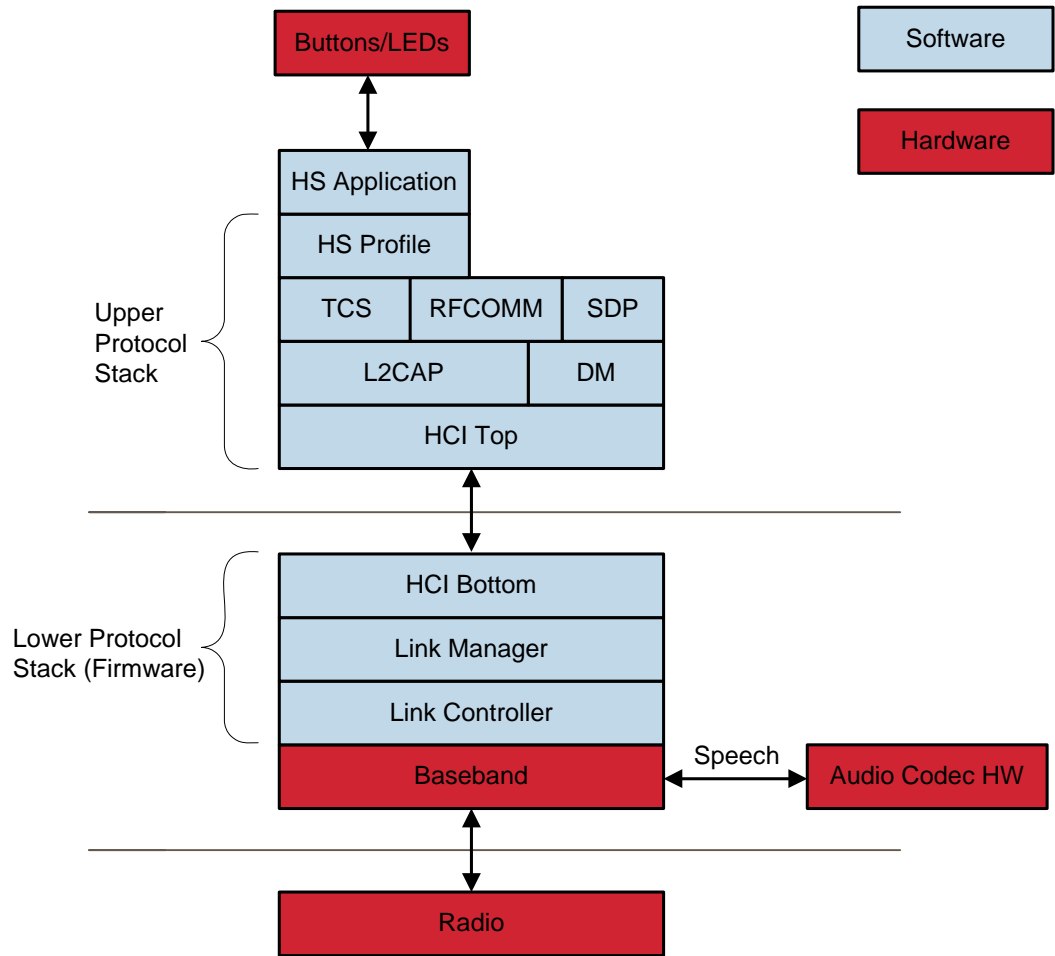
## SPEECH QUALITY ONLY

### *Bluetooth HW/SW Architecture*

A complete Bluetooth solution consists of multiple components:

- The radio transmits and receives Bluetooth signals over the air. It typically consists of analog (RF) and digital hardware components (mixed-signal technology).
- The baseband provides link controlling and link managing (Link Controller/Manager). The functionality is partitioned between digital hardware and software running on a microcontroller. This software is called "Lower Protocol Stack," which is actually the device firmware. The Bluetooth specification defines an interface called the "Host Control Interface" (HCI) to access the lower layers, as shown in Figure 1.

- The upper layer software that accesses the lower layers through HCI is called the “Upper Protocol Stack.” It typically runs on a separate microprocessor and includes the higher level components of the Bluetooth Protocol Stack. The processor also runs the application, which interfaces to the user and to the Upper Protocol Stack APIs.



**Figure 1. Bluetooth Protocol Stack with Headset Profile**

An alternative architecture is a one-processor solution. In this case, the entire Bluetooth Protocol Stack runs on a single processor, and the HCI is not visible to the outside world. This solution, however, imposes challenges to fulfill the stringent real-time requirements of the Link Controller/Manager, especially when the application requires significant processing power.

Also, the firmware needs to address the double memory range, necessitating a higher performance microprocessor. Changes in your application require a full rebuild and test. In addition, this architecture typically imposes inefficiency in terms of power consumption due to the larger chip area and higher clock frequency required.

#### *SCO and eSCO*

The Headset Profile uses an SCO (Synchronous Connection Oriented) link. Bluetooth defines this special packet type for telephony services. Due to latency requirements, this packet type does not use CRC or retransmission and is therefore called “unreliable.” Speech data is typically not transmitted over the HCI but is instead managed directly by the baseband.

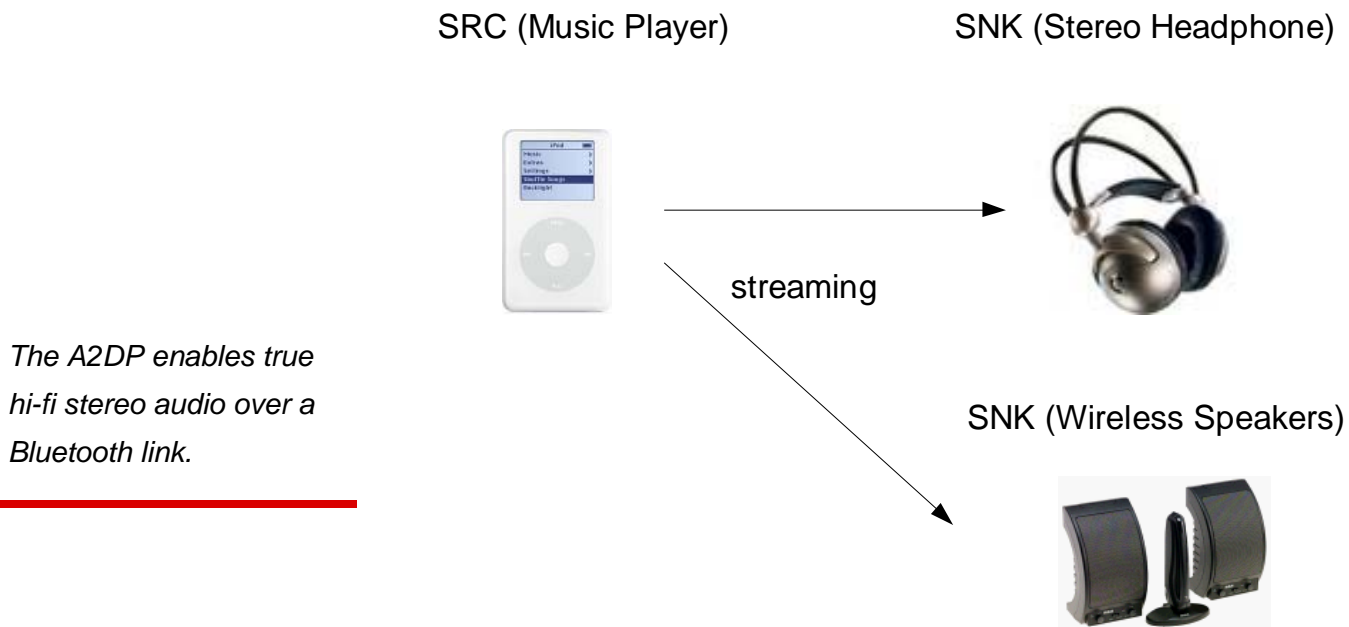
Bluetooth Specification 1.2 introduces eSCO, which provides “reliable” connections by using CRC and retransmissions. It also features more bandwidth for speech than the 64-kbps SCO (wideband speech). Unfortunately, there is a large set of parameters that the two devices need to agree on, making a fallback to the old SCO link likely. In addition to this limited interoperability, there is also no efficient speech codec standardized for bit rates higher than 64 kbps, which makes widespread use of eSCO more unlikely.

## **HI-FI STEREO QUALITY**

Since eSCO is not a suitable candidate for high-quality audio streaming, system designers need to rely on ACL (Asynchronous Connectionless) links. Packet-wise transmission with CRC and retransmission guarantees the reliability needed.

The new A2DP defines different codecs to use for efficient transmission of audio data. The SBC codec is mandatory and guarantees interoperability between different Bluetooth devices. Other codecs like MP3 and AAC are optional.

The source (SRC) or the music player (for example, iPod) connects to one or more sinks (SNKs), which are basically the stereo headsets. After successfully establishing a connection, the SRC starts streaming ACL data containing compressed audio. The SNK decompresses the audio data and outputs it to the speakers.



*The A2DP enables true hi-fi stereo audio over a Bluetooth link.*

**Figure 2. A2DP Enabled iPod Streaming to Wireless Speakers and Stereo Headset**

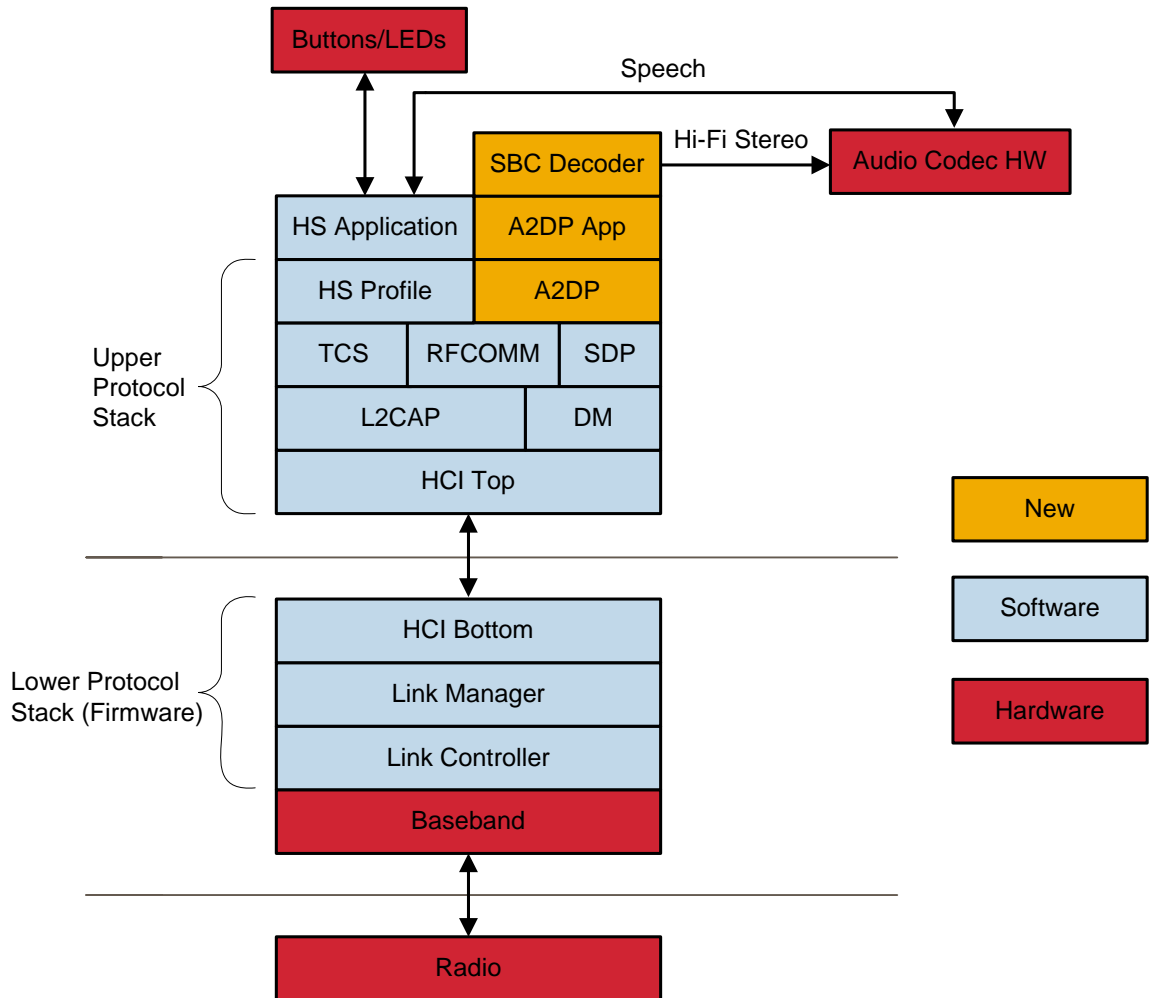
## CHALLENGES

At first sight, implementing the A2DP seems to be straightforward; however, it actually poses several problems you need to deal with.

The most important change in the architecture compared to the Headset Profile is that you need a decoder for SBC packets. The application running on the SNK has to deal with the decoding directly, whereas in the Headset Profile this burden is typically delegated to the baseband. This implies that ACL packets containing SBC audio data have to be transmitted over the HCI, which imposes significantly larger bandwidth requirements than before since SBC bit rates may go up to 510 kbps.

Decoding the SBC packets requires a dedicated DSP, thus another microprocessor in the Bluetooth system. In total, you come up with three different microprocessors, likely all with different toolchains for software development. Engineers need to get familiar with all of them, prolonging development cycles and raising project costs.

It is also difficult to ensure proper inter-communication between the components. That is, managing the data flow through the buffers between the computing components becomes a tricky task. Latency and synchronization are the key issues that arise.



**Figure 3. Bluetooth Protocol Stack with A2DP and Headset Profile**

The brand or type of DSP you need is dictated by the availability of an SBC decoder software implementation. Special care needs to be taken since the software has to be “qualified,” which means it needs to pass the conformance tests imposed by the Bluetooth SIG (Special Interest Group). Also, the implementation needs to be efficient. Otherwise, the DSP drains too much battery power and requires too much memory, which increases total system cost. Finding a suitable DSP is most critical especially when it needs to fit into the overall architecture.

## FROM SPECIFICATION TO IMPLEMENTATION

Several architectures are possible, each with several benefits and drawbacks. An architecture where the HCI is accessible is called “hosted”.

### *One Microprocessor*

In this scenario, the entire Bluetooth Protocol Stack and the SBC decoder run on a single processor. As an advantage, you only need to have one toolchain for development. Less time needs to be spent on ensuring proper inter-communication between other processors. On the other hand, testing your system becomes difficult, since there are no physical interfaces visible to the outside world. Also, the stringent real-time requirements put a high burden on the processor, especially since it also runs the SBC decoding. The higher clock rates required imply higher power consumption. To make it worse, there are often times when parts of the system are not active, like when the A2DP is not in use. Also, an efficient SBC decoder is likely not available, thus making this architecture unrealistic.

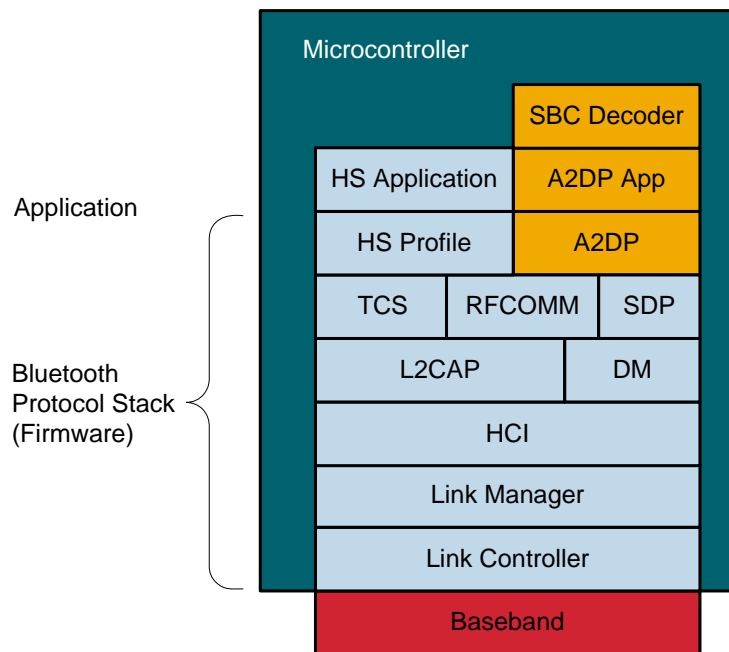
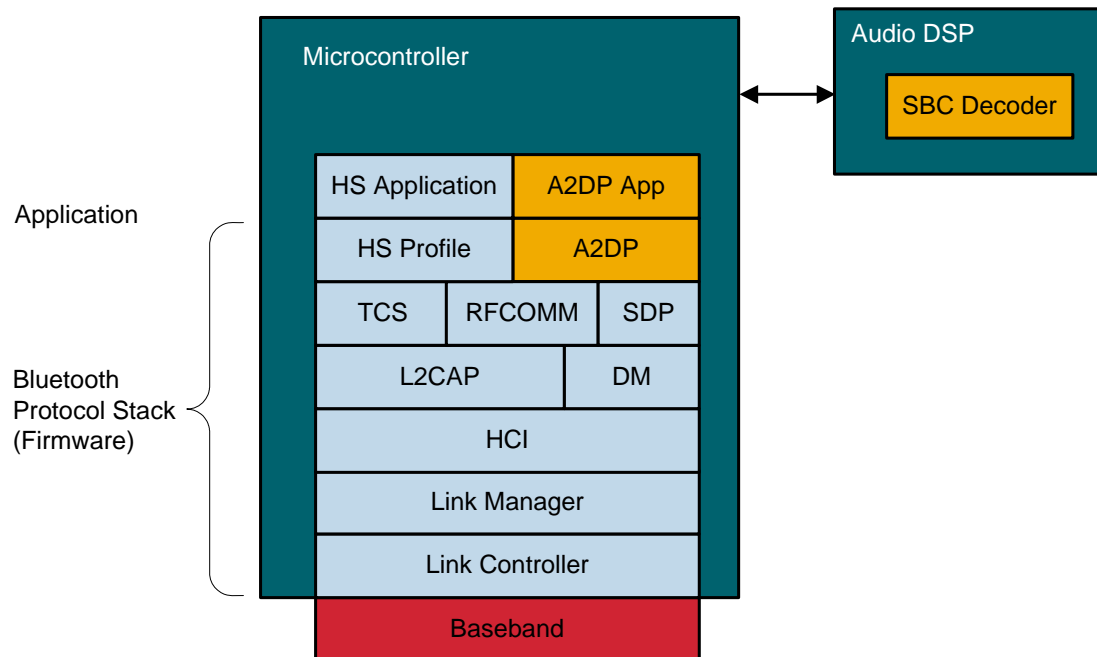


Figure 4. One Processor Solution



### *One Microprocessor and One DSP*

Here, one processor runs the entire Bluetooth Protocol Stack and a DSP runs the SBC decoder. This architecture is feasible but is difficult to test. Also, running Upper and Lower Protocol Stack on a single processor makes it difficult to maintain the real-time requirements, thus unnecessarily driving up clock frequency and increasing power consumption.

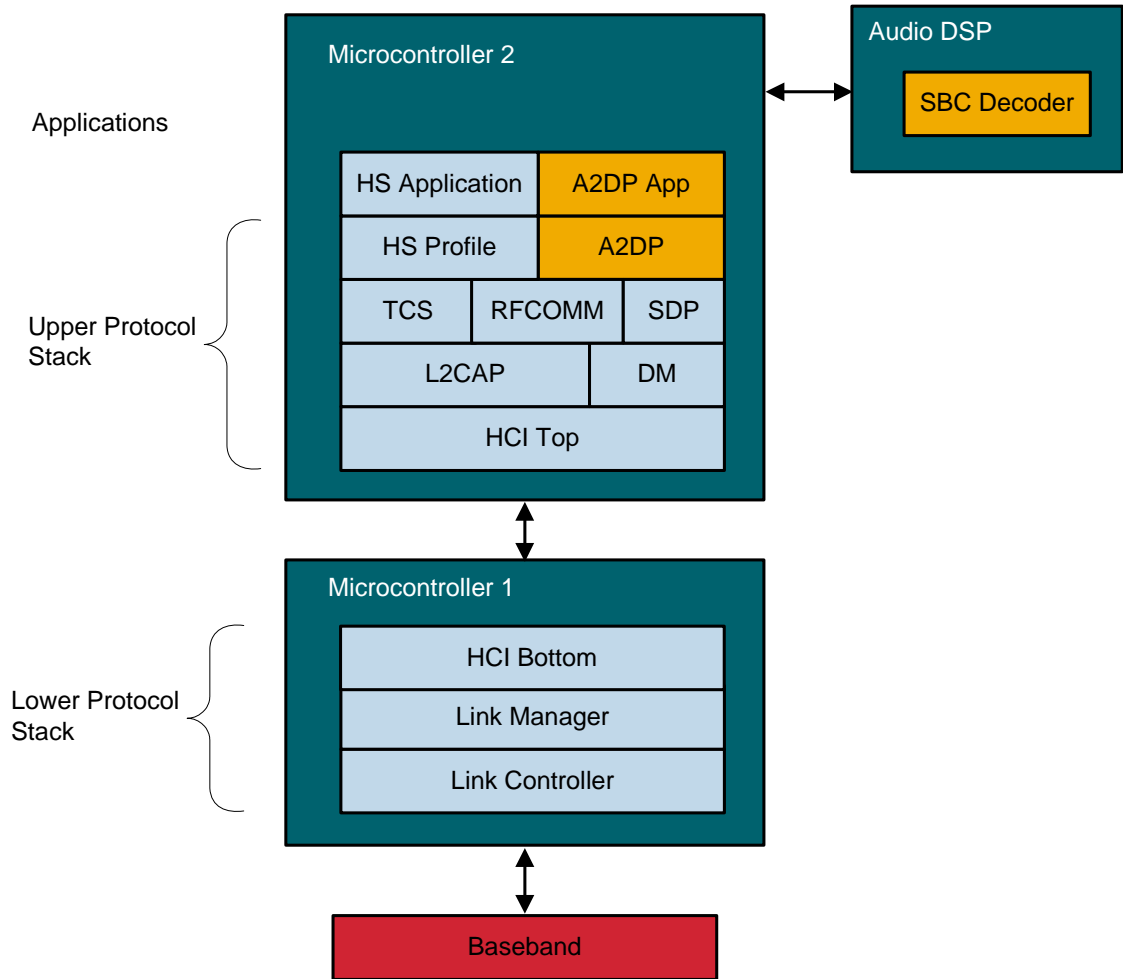


**Figure 5. One Processor Plus DSP**

### *Two Microprocessors Plus One DSP*

A three-processor solution (two microcontrollers and a DSP) presents clearly defined interfaces, easing integration and test—and thus time to market. Also, a small memory footprint together with low MIPS requirements lead to low power requirements and low cost. Fewer interrupts per processor ameliorate interrupt latency, allowing a lower clock frequency.

On the downside, you need to deal with toolchains for three different targets. Also, the tasks of integration and exchanging data between the different processors increase the engineering effort. Utilizing three different processors increases total system cost.



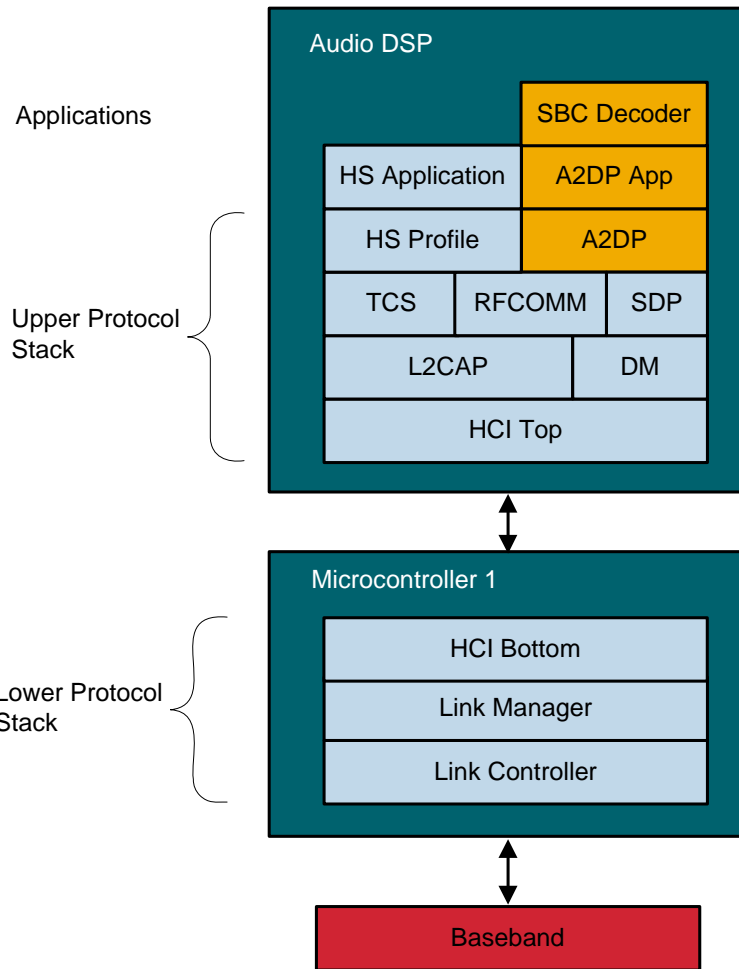
**Figure 6. Two Processors Plus DSP (Hosted)**

*IPextreme’s Two-Processor Solution*

IPextreme developed a unique solution that combines the advantages of the different architectures but avoids their drawbacks.

We use a dedicated processor to run the firmware and use HCI to interface the upper components of the stack. The Upper Protocol Stack and the A2DP SNK application both run on a single DSP, which also decodes the SBC packets if A2DP is selected. By managing to run control-type software (the Upper Protocol Stack) on a dedicated audio DSP in an efficient manner, we are able to obtain significant savings in chip area, resulting in lower power consumption and lower total system cost. Also, integration effort is lower since you only need to work with two toolchains.

*Running the Upper Stack and SBC Decoder on a single DSP optimizes the processor usage while keeping the Bluetooth Lower Stack real-time processing isolated in a dedicated microcontroller .*



**Figure 7. One Processor Plus DSP (Hosted)**

**PUTTING IT ALL TOGETHER**

*Components used for Upgrade from Headset Profile*

At IPextreme, we used IP from our partners to build a prototype.

The NXP (formerly Philips) CoolFlux DSP is a dedicated ultra-low-power DSP for audio applications. As a very unique feature, it uses a retargetable toolchain from Target Compiler Technologies. This breakthrough technology allows you to instantly generate a new version of the toolchain after making changes to your processor architecture. By default, the toolchain exploits all parallelism in the processor and produces highly optimized machine code from C-code.

For example, if you wish to extend the addressable memory, first you apply the changes to the CoolFlux DSP RTL. Second, in the toolchain, you only need to update the processor model to reflect the changes you did in the RTL. The toolchain (including C-compiler) is ready immediately after a change in the processor model.

A CoolFlux DSP optimized version of the SBC decoder from Philips fulfills the requirements of efficiency and conformance to the Bluetooth Specification.

The Bluetooth Protocol Stack from Cambridge Consultants, the market leader in Bluetooth technology, ensures interoperability.

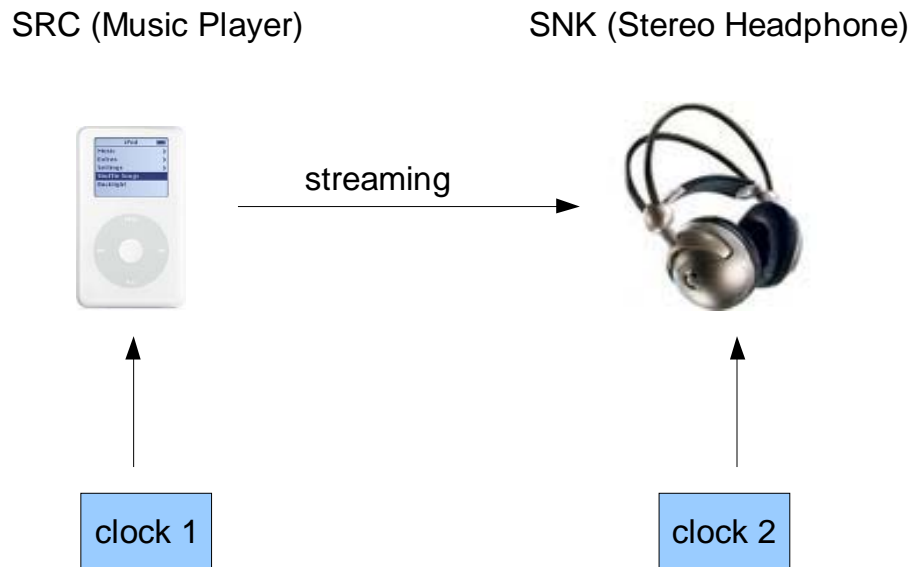
#### *Tasks for Integration*

At IPextreme, we ported the Bluetooth Protocol Stack to the CoolFlux DSP platform. We then built an A2DP SNK application that interfaces to the SBC decoder. In addition, we developed drivers to interface the audio codec hardware that connects to the speakers and the UART driver to ensure proper communication over HCI.

In a later step, we optimized the memory buffers to allow proper communication and data flow between the Upper Protocol Stack and the SBC decoder.

As the Bluetooth Specification does not give a solution on how to provide intra-stream synchronization, we developed an algorithm that renders the PCM audio to the audio codec hardware. Not having this would result in skips or dropouts of audio, when the SRC either sends too much or too little data. As the Bluetooth Specification does not define any mechanisms to synchronize the clocks of the SRC and SNK, a synchronization algorithm is required.

The prototype was proven to have interoperability with multiple Bluetooth Protocol Stack providers.



**Figure 8. System with Two Independent Clocks**

## BENEFITS

Putting together a product that involves multiple components of complex hardware and software IP can be done easily. An IP solution allows you to customize the hardware to your specific needs if you wish to. At the same time, a retargetable toolchain allows you to keep the development tools in sync with your modifications to the RTL. Help from the toolchain provider is not necessarily required.

Ease of integration together with low cost are the major factor for companies wanting to integrate Bluetooth stereo into their products.

## SUMMARY

Successful integration of numerous complex IP components is feasible today. High demand for Bluetooth stereo today highlights the need for an IP solution with a quick time-to-market. IPextreme is committed to help you succeed in Bluetooth stereo technology. A demo called "CoolBlue" (CoolFlux + Bluetooth) is available today and has already been demonstrated at numerous shows.

*Almansor Kerroum is a senior design engineer with IPextreme. He received his MSEE from University of Kaiserslautern, Germany. He then spent three years working on wireless algorithm development before joining IPextreme. He can be reached at [almansor.kerroum@ip-extreme.com](mailto:almansor.kerroum@ip-extreme.com).*



[www.ip-extreme.com](http://www.ip-extreme.com)

---

**IPextreme, Inc.**

307 Orchard City Drive

M.S. 202

Campbell, CA 95008

800-289-6412 (toll-free)

408-608-0421 (fax)

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND. INFORMATION IN THIS DOCUMENT IS SUBJECT TO CHANGE WITHOUT NOTICE.

© Copyright 2006, IPextreme. All rights reserved. IPextreme and the IPextreme logo are trademarks of IPextreme, Inc. All other trademarks are the property of their respective owners.

---