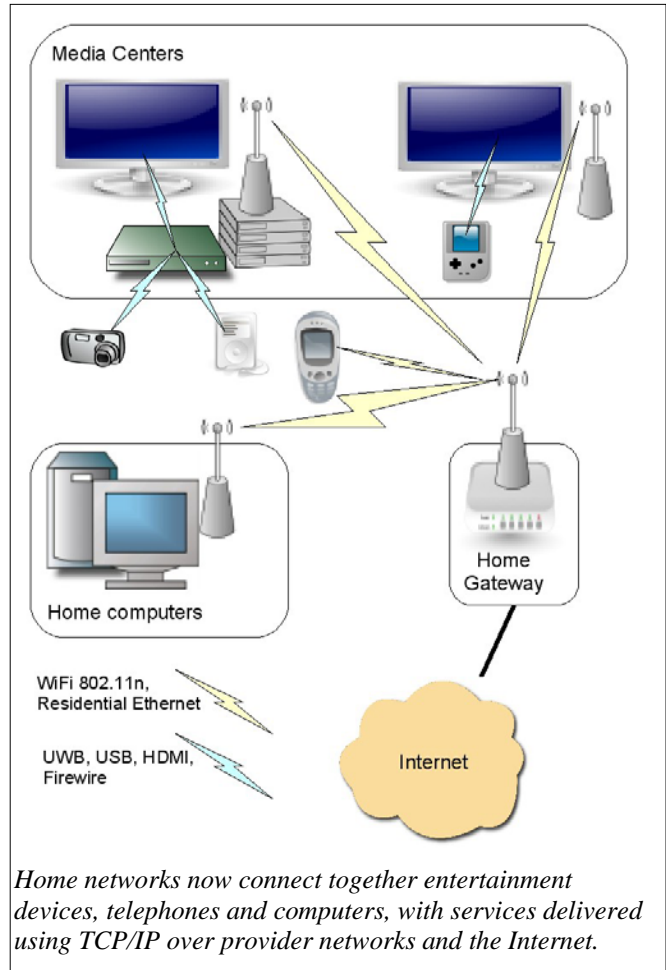


An Overview of Digital Rights Management for System-on-Chip Architects

Introduction

Artists and studio understand the power of digital distribution of content with the dramatic success of personal entertainment devices such as the Apple iPod. Anxious to embrace this revolutionary and highly effective channel, content owners and service providers are embracing digital distribution using a wide variety of digital rights management (DRM) technologies such as Apple's proprietary iTunes technology, Windows Microsoft DRM, the Open Mobile Alliance (OMA) standard or the Digital Transmission Licensing Authority DTCP standard. This wide variety of designs has made it challenging for IC designers. Indeed the correct solution to this challenge is a multi-DRM design implementation that spans as many DRM schemes as possible to ensure that each new cell phone or set-top box chip can access as broad a range of content as possible.

Adding to the complexity is the fact that consumers now demand that devices that formerly did not communicate together now routinely share data so that consumers may enjoy the content they licensed anywhere in their home and when they're traveling regardless of where the content originated or is stored. This unprecedented mobility means that consumer electronics devices must now be capable of dealing with different DRM technologies from such diverse originators as cable and satellite TV, mobile phone operators, multimedia computers and content producers. This whitepaper discusses the issues faced by SoC architects responding to this challenge.



Background

DRM systems are a basic component of modern digital media devices. Content providers from traditional media outlets such as the audio and video recording industry, television broadcasters and movie studios deliver programming using a multitude of different DRM models, both standards-based and proprietary.

The number of DRM models in a typical household is increasing rapidly as previously disparate networks and devices are connected together to share data. Consumers are consolidating their media collections in on-line libraries stored in media PCs or network-attached disk storage.

Proprietary and stand-alone systems to play and distribute media are giving way to home entertainment networks that connect personal computers, playback devices like DVD and MP3 players, Internet-enabled game consoles, HDTV monitors and televisions, and next generation set-top boxes (STB) and personal video recorders (PVRs) using high speed Ethernet and WiFi networks. Next generation broadband networks provide converged services like VOIP, digital TV and Internet access, all delivered on an Internet Protocol data network.

To put this into context, one must consider how we got to where we are today. Two major classes of devices represent the historical basis for modern DRM systems: cable and satellite television subscription services; and wireless phone and related services. After initial disastrous experiments with “security by obscurity” that made it easy for unauthorized users to access programming, television subscription service providers developed Conditional Access (CA) systems that embody most of the same elements of modern DRM. CA systems use several distinct security services including subscriber authentication and subsequent key distribution to enable authorized subscribers to enjoy encrypted streaming content. Most systems like this provide limited protection against unauthorized deciphering of encrypted content and are suitable only to discourage casual theft.

Mobile wireless telephone systems underwent a similar early experience as operators came to understand that it was impossible to prevent increasing losses to unpaid, unauthorized access to their networks by doing little more than embedding password access codes in devices accessing their networks. What resulted was the underground “clone-phone” industry in which many phones would be programmed with the same valid authentication codes and sold to unauthorized users. Embedded passwords have given way to Subscriber Identity Module (SIM) cards based on smartcard technology using strong cryptographic

techniques to authenticate devices to the network. (While the techniques are generally secure, weak implementations have meant that the security of the system is less than originally intended.) Some cellular networks provide encryption of the content (the phone call), although the security has been relatively low.

At about the same time, music and movie studios and other program content providers started to recognize that emerging digital technologies offered compelling reasons for consumers to adopt them. They also saw in that the threat of perfect digital copies of programming distributed at low cost and high volume over the Internet. This led content providers to find their own methods to protect their content from unlimited reproduction and redistribution. Different DRM technologies started to appear in the form of CSS, Windows DRM, DTCP, OMA and others. In general, these are incompatible with each other and require different underlying functions to implement.

DRM applications

In the early days of DRM, most systems had to contend with one, or perhaps two, different DRMs. Today manufactures of integrated circuits and systems find themselves providing multiple DRM implementations to translate between one DRM and another. For example, modern set-top boxes take in content from DVDs (CSS), cable systems (CA) and other sources, store it for playback on hard disk, then redistribute that content to DTCP-enabled displays around the house via Firewire or WiFi networks. Each stage in this process requires cryptography at HDTV data rates. This places incredible processing load on embedded processors in these device. A single 25 megabit per second video stream can require more than a billion instructions per second of CPU processing on an embedded processor, and frequently several streams are being processed simultaneously. For example, modern PVRs allow simultaneous

playback of previously recorded programs while recording one or more programs to hard disk for later playback.

Consumers now enjoy very high quality surround-sound audio to complement the video in home theater systems. These systems are just undergoing the digital revolution, promising easy distribution of home theater content around the house without the need for expensive custom installations. To respond to this need, IEEE has just embarked on development of the Residential Ethernet (ResE) standard. This will provide the highly accurate synchronization of signals needed to provide high resolution imaging of the audio field to produce enjoyable surround sound. Somewhat surprisingly this places extra demands on DRM systems which must now provide low-latency highly predictable and synchronized processing of multiple channels of data. Just as in data networking today, the ResE initiative can be expected to garner corresponding activity from wireless networking technologies such as WiFi and ultrawideband (UWB) to produce similar capabilities. Outstanding hardware implementations to produce the required predictability and latency characteristics will be required to service these applications as software only implementations cannot support these performance requirements..

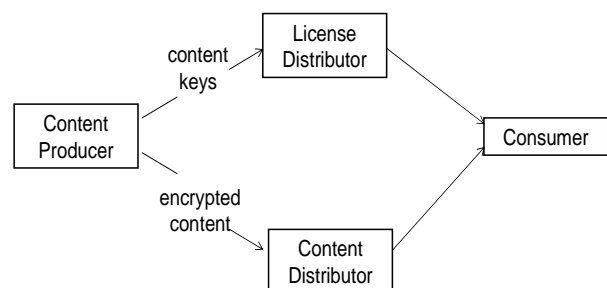
When the bandwidth and latency characteristics of DRMs were not terribly demanding, manufacturers provided software implementations. This was attractive as it allowed easy upgrades for their systems. The most efficient systems allow end-users to download and install upgrades over the same network that content is delivered on. This minimizes support and distribution costs for manufacturers, but creates challenges for both system manufacturers and content providers. Always-on Internet-connected devices chock full of movies and music are obvious targets of interest for hackers and malicious software. Of particular interest are the stored content and corresponding content decryption keys, as well as the software that

operates the device itself. While consumers are accustomed to thinking of STBs and recording devices as hardware, hackers see them for what they are: a potential server capable of distributing gigabytes of content and other data over high-speed network connections.

For this reason, DRM licensing bodies often place severe requirements on device manufacturers and implementers to protect against leaks of protected content and encryption keys, even in the face of a hostile determined attack. Licenses for proprietary DRM technologies often include provisions for large financial liabilities for allowing such leaks to occur. Manufacturers of consumer class devices must now adopt techniques and processes traditionally reserved for applications like Internet commerce (websites such as banks) to protect against their devices becoming open conduits for illegal activity.

Elements of DRM

Regardless of the particular DRM implementation, all DRMs share the same basic set of functional elements. Content producers prepare their content for distribution using tools supplied by the DRM manufacturer. Content files are encrypted using symmetric encryption (typically AES). In general, a single encryption key is used uniquely for a particular title.

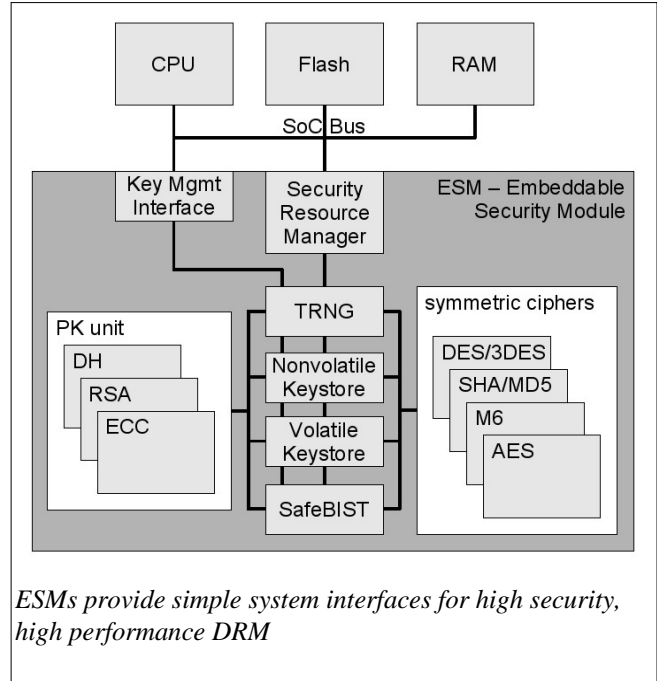


Once protected, the content file can be moved to the distribution channel. Any number of possible distribution models are possible including privately owned server networks, traditional retail distribution on hard media such as CDs and DVDs, on-line distribution on private networks such as cablevision, or even posting on public web servers for distribution on the Internet. The corresponding content encryption key is itself protected and moved to the license management system of the DRM.

The license management system includes features such as subscription management, interfaces to billing systems, and importantly for on-line services a license distribution function. A consumer wishing to access protected content contacts the license server and provides proof of legitimate rights to access the content in the form of a cryptographic certificate. If the license distributor can authenticate the right of the consumer to access the content, it returns a copy of the content key encrypted using a public key in the consumer's authentication credentials. The content key can only be decrypted using the corresponding private key.

Distribution on hard media is often dealt with by providing the license keys in files stored in private regions of the media. These typically use non-standard proprietary schemes to implement. The private regions are designed such that they are not reproduced using consumer-grade media copying equipment (e.g. DVD). These schemes often provide only very limited protection against casual attempts to access or reproduce material without legitimate license.

Early on, it was common to use pre-shared keys (key encryption keys) to protect the content keys. Any compromise of the pre-shared keys resulted in compromise of the entire system. Modern DRMs use public key cryptography in combination with key distribution protocols to protection the content keys. Compromise of a particular content



protection key, while effectively making the content available to everyone, does not compromise the entire system. Similarly, compromising the private key of an individual device does not affect the security of other devices that do not share the same key. This makes the overall system much more robust to failures of individual parts of the network.

There are a few public key cryptographic algorithms used in DRMs, and of those the two most commercially important are the Elliptic Curve Cryptography (ECC) and RSA algorithms. RSA is the older and more established of the two having been invented more than 20 years ago and has been a workhorse for years. ECC enjoys several advantages over RSA. At a given key size, ECC is significantly more secure than RSA in face of an exhaustive attack on the keyspace. In practical terms ECC systems can use much smaller keys than RSA to achieve the same level of security. (When comparing crypto-systems, it is essential to do so at equivalent levels of security. The discussion below makes this assumption implicitly.) This in turn produces several benefits:

- ECC can produce acceptably fast results in software for many applications. RSA often requires dedicated hardware to achieve acceptable performance.
- Owing to its more efficient computational structure ECC consumes less power than RSA. This can be extremely important in battery-powered applications.
- When storing keys in expensive nonvolatile memory, ECC requires only about 10% to 20% of the space of equivalent RSA keys.
- Hardware-assisted ECC scales to an order of magnitude or more operations per second than RSA.

As the target security level rises, the advantages of ECC over RSA also increase due to ECC's exponential increase in strength with increasing key size. This makes the future clearly tilt in favor of ECC. However, legacy systems built on RSA will continue to exist for years and SoC's must be designed to deal with both ECC and RSA in multiple DRM environments.

Hardware design implications

One of the key jobs of the SoC architect is partitioning functionality between hardware and software. Beyond the usual performance and cost trade-offs, DRM-enabled systems also have security as a design objective. A good system design will deal with this explicitly, weighing cost and complexity against the ability of the system to withstand attacks from defined threats. These systems share many of the characteristics of Hardware Security Modules (HSMs), and it is instructive to examine the techniques of HSM design to understand them better. To this end, resources such as NIST FIPS 140-2 and related documents may be helpful.

Elliptic Technologies's Embeddable Security Module (ESM) products have encapsulated these techniques to simplify the task for SoC architects. ESMs provide simple system interfaces and resources to facilitate a range of security functions including

- secure nonvolatile key storage
- validation and self-test of cryptographic resources at system start-up and ongoing
- true random number generation at entropy levels consistent with the overall system security level
- high performance symmetric and asymmetric cryptographic operation implementations
- separation of key management from key usage
- binding of cryptographic keys to their designated use to prevent bootstrapping attacks to attempt recovery of the keys

Summary

Elliptic Technologies has extensive experience in multi-DRM solution including IP cores and associated software. The company has implemented designs for mobile and set top box applications encompassing Windows DRM, OMA 2.0, DTCP, IPsec and proprietary DRM. With its software partner Certicom, Elliptic can offer solutions that span Technologies IP all the way up to content servers and license distribution. For more information on Elliptic products and services please contact us at:



Embedded security you can trust

White Paper Digital Rights Management

Elliptic Technologies Inc.
Steacie Dr., Suite 201
Kanata, ON
K2K 2A9
Phone: 613 254-5456
info@elliptictech.com