

Software Defined Radio

**Baseband Technologies for 3G Handsets
and Basestations**

Edited by

Walter H.W. Tuttlebee

Virtual Centre of Excellence in Mobile and Personal Communications (Mobile VCE)



John Wiley & Sons, Ltd

WILEY SERIES IN SOFTWARE RADIO

Series Editor: Dr Walter Tuttlebee, Mobile VCE, UK

The Wiley Series in Software Radio aims to present an up-to-date and in-depth picture of the technologies, potential implementations and applications of software radio. Books in the series reflect the strong and growing interest in this subject. The series is intended to appeal to a global industrial audience within the mobile and personal telecommunications industry, related industries such as broadcasting, satellite communications and wired telecommunications, researchers in academia and industry, and senior undergraduate and postgraduate students in computer science and electronic engineering.

Mitola: Software Radio Architecture: Object-Oriented Approaches to Wireless Systems Engineering, 0471384925, 568 Pages, October 2000

Mitola and Zvonar (Editors): Software Radio Technologies: Selected Readings: 0780360222, 496 pages, May 2001


Tuttlebee: Software Defined Radio: Origins, Drivers and International Perspectives, 0470844647, £65, 350 pages, 2002

Tuttlebee: Software Defined Radio: Enabling Technologies, 0470843187, £65, 304 pages, 2002

Dillinger, Madani and Alonistioti (Editors): Software Defined Radio: Architectures, Systems and Functions, 0470851643, £85, 456 pages, 2002

Tuttlebee: Software Defined Radio: Baseband Technology for 3G Handsets and Basestations, 0470867701, £75, 275 pages

Software Defined Radio



WILEY SERIES IN SOFTWARE RADIO

Series Editor: Dr Walter Tuttlebee, Mobile VCE, UK

The Wiley Series in Software Radio aims to present an up-to-date and in-depth picture of the technologies, potential implementations and applications of software radio. Books in the series reflect the strong and growing interest in this subject. The series is intended to appeal to a global industrial audience within the mobile and personal telecommunications industry, related industries such as broadcasting, satellite communications and wired telecommunications, researchers in academia and industry, and senior undergraduate and postgraduate students in computer science and electronic engineering.

Mitola: Software Radio Architecture: Object-Oriented Approaches to Wireless Systems Engineering, 0471384925, 568 Pages, October 2000

Mitola and Zvonar (Editors): Software Radio Technologies: Selected Readings: 0780360222, 496 pages, May 2001

Tuttlebee: Software Defined Radio: Origins, Drivers and International Perspectives, 0470844647, £65, 350 pages, 2002

Tuttlebee: Software Defined Radio: Enabling Technologies, 0470843187, £65, 304 pages, 2002

Dillinger, Madani and Alonistioti (Editors): Software Defined Radio: Architectures, Systems and Functions, 0470851643, £85, 456 pages, 2002

Tuttlebee: Software Defined Radio: Baseband Technology for 3G Handsets and Basestations, 0470867701, £75, 275 pages

Software Defined Radio

**Baseband Technologies for 3G Handsets
and Basestations**

Edited by

Walter H.W. Tuttlebee

Virtual Centre of Excellence in Mobile and Personal Communications (Mobile VCE)



John Wiley & Sons, Ltd

Copyright © 2004

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk
Visit our Home Page on www.wileyeurope.com or www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770571.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Trademarks of all companies referred to in this book are acknowledged.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 33 Park Road, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

Library of Congress Cataloging-in-Publication Data

A catalog record for this book is available from the Library of Congress

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 0-470-86770-1

Typeset in 10/12pt Times by Integra Software Services Pvt. Ltd, India

Printed and bound in Great Britain by TJ International Ltd, Padstow, Cornwall

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

Contents

List of Contributors	xv
Foreword <i>Stephen M. Blust</i>	xix
Abbreviations	xxi
Biographies	xxix
Introduction <i>Walter Tuttlebee</i>	1
Part I: Requirements	5
1 SDR Baseband Requirements and Directions to Solutions	7
<i>Mark Cummings</i>	
1.1. Baseband Technology and the Emergence of SDR	7
1.1.1. <i>Wireless Architectures</i>	7
1.1.2. <i>The Impact of Digital Technology</i>	9
1.1.3. <i>Growth of the Wireless Markets and Multiple Standards</i>	10
1.2. Evolution of SDR Baseband Requirements	12
1.2.1. <i>System Optimization</i>	12
1.2.2. <i>Market Drivers of SDR</i>	12
1.2.3. <i>Technology Drivers of SDR</i>	13
1.2.4. <i>Requirements by Application Environment</i>	14
1.3. Mobile End-user Systems	15
1.3.1. <i>Simple Handsets</i>	15
1.3.2. <i>Converged Devices</i>	18
1.3.3. <i>Notebooks</i>	19
1.3.4. <i>Vehicular Systems</i>	20
1.4. Fixed End-User Systems	22
1.4.1. <i>Infrastructure System Requirements</i>	22
1.4.2. <i>Basestation Requirements</i>	23

100

100

100

100

1

SDR Baseband Requirements and Directions to Solutions

Mark Cummings

RFco, Inc.

This chapter puts the development of SDR baseband technology into functional and historical perspective, then examines the evolving spectrum of requirements from mobile devices such as simple handsets, through complex mobile devices, access points, micro cellular basestations to wide area basestations. Both internal functional requirements and external interface requirements are considered. The requirement challenges are then discussed in the context of currently available (and emerging) technologies for implementing SDR baseband processing subsystems. The chapter concludes with a consideration of likely successful implementation directions.

1.1. Baseband Technology and the Emergence of SDR

The primary intent of SDR is to allow a wireless device to change functionality by switching in different software stacks (or software-like objects) that can be stored locally or downloaded over the air [1].

1.1.1. *Wireless Architectures*

At the highest level, all wireless systems can be decomposed into four subsystems (see Figure 1.1). The baseband subsystem sits between the RF front-end and the controller. It is responsible for end-user data encoding/decoding and signal modulation/demodulation. The encoding/decoding function can be classified as a low speed signal-processing task and the modulation/demodulation task can be classified as a high speed signal-processing task.

By the end of the Second World War, wireless communication systems had yet to evolve to the functional form shown in Figure 1.1. Rather they consisted of discrete analog electrical




Figure 1.1 A top level view of wireless systems architecture




Figure 1.2 Evolution of wireless systems from discrete analog to software-driven hardware

components that performed the following functions (also listed at the top of Figure 1.2, from right to left):

- Human interface,
- Local control,
- Protocol stack,
- Low speed signal processing,
- High speed signal processing,
- RF front-end,
- Antenna.

Examples of the tasks performed by each of these functional units, in a cellular handset include:

- Human interface – controlling the speaker, microphone, keyboard and display;
- Local control – managing the handset hardware and software resources;
- Protocol stack – managing such functions as call set-up and maintenance;
- Low speed signal processing – user data coding, such as interleaving;
- High speed signal processing – modulation such as FM, QPSK, GMSK;
- RF front-end – up/down conversion, amplification and filtering;
- Antenna – interface to the ether.

In these early systems, each function was implemented with discrete analog technology. This resulted in relatively large, expensive, high power-consuming systems, which were difficult to design, manufacture and manage/maintain in the field. The desire to reduce the cost, size and power consumption, and to make devices easier to manage in the field while taking full advantage of technology improvements wherever possible, has driven the technology evolution path we are still on today.

1.1.2. The Impact of Digital Technology

As digital technology arrived and entered the beginning of its period of rapid evolution, a pattern developed. First a function previously performed in discrete analog circuitry was implemented with discrete digital components. Then, as the semiconductor technology developed, these discrete components were integrated into larger and larger units and the specialized, discrete, hard-wired digital logic was replaced with general purpose logic that changed its function with changes in software. This combination of general purpose logic plus software on a single semiconductor came to be called a microprocessor. Several types of microprocessor evolved, each optimized for different functionality. This process is graphically portrayed in Figure 1.2.

1.1.2.1. Early Applications in Radios

The first functions to be aggregated through this process were the human interface, local control, and protocol stack functions. With the appearance of the microprocessor, the discrete logic components that replaced the analog components were, in turn, replaced with a specialized microprocessor, called a microcontroller, and software. Then the low speed signal processing function went from analog discrete components through digital logic to be implemented as another specialized processing engine. Special mathematical functionality (such as multiply accumulate, MAC) was added to microprocessors to create digital signal processors (DSPs) enabling the low speed signal-processing functions to be converted from discrete digital logic to software-programmed DSPs .

1.1.2.2. The Dream of 'Pure' SDR

Next, the high speed signal-processing analog discrete components were replaced with digital logic components. The expectation was that the same process would continue and that high speed signal processing would also soon be implemented by some kind of microprocessor and software. The dream was that this process would continue to the point where a receiver subsystem would consist of an A/D converter at the antenna and everything else would be done by software. However, another factor was by now beginning to come into play – the increasing complexity of the radio systems themselves.

1.1.2.3. The Growing Complexity of Cellular Radio – DSPs and ASICs

Early cellular standards such as AMPS required a baseband bandwidth of 30kHz. As low speed signal processing was moving into a DSP implementation, second generation cellular standards such as GSM emerged, with a baseband bandwidth in the 200 to 300kHz range.

Other second generation standards, such as IS-95 (CDMA/AMPS dual mode air interface standard, AIS) with an even larger baseband bandwidth of 1.24 MHz, were also coming to the market.

Although monolithic, fixed, single-stream instruction set processor DSPs could theoretically handle the speed of processing required for high speed signal processing of second generation AISs, practical limitations created an impenetrable barrier for the traditional DSP, with its single-stream instruction set processor. This architecture has a single processing unit which must load data (single sample of signal data) and a single instruction, process the data, write the result to a buffer and start over again. The Nyquist theorem requires that a signal be sampled at more than twice its rate of change (in practice 2.5 times its bandwidth in Hertz) in order to preserve the data in the signal. This means that a 1.24-MHz signal must be sampled at approximately 3.25 MHz. In order to avoid quantization error, single samples are typically represented by 13 or more bits in cellular systems. Minimum high speed signal processing for first generation and early second generation AISs require approximately 100 instructions per sample. Assuming that the processor has a 16-bit word size, this means that a DSP attempting to do high speed signal processing for a simple handset would have to operate at a clock speed in excess of 325 MHz. In practice, because of bus delays, the need to write to buffers, etc., the required clock speed turns out to be in the GHz range. Early DSPs could not run at these clock speeds.

As DSP development progressed, it became clear that, since power consumption varied directly with processor speed, it would not be practical to operate a DSP at these clock rates for battery powered applications. It also turned out that, even for systems run on utility-provided power, the nature of their requirements plus issues of size, heat dissipation, etc., would make pure DSP solutions for high speed signal processing impractical.

Discrete digital component implementations avoided the power consumption problem by implementing each of the 100 instructions in 100 discrete digital logic circuits arranged in a bucket brigade. In this way, each circuit could run at the (much lower) clock rate of the sampled and quantized signal (such as 3.25 MHz), dramatically lowering the power consumption. Further power/speed improvements can be made by having each circuit optimized for a single function, thereby avoiding the inefficiencies inherent in a general purpose processor which has to be optimized for a wide range of functions.

So, here, the previous pattern appeared to break. Discrete digital logic components could not be replaced by a microprocessor with software. However, they could be combined into a single chip through the aggregation and integration of the discrete logic components into a single large scale integrated circuit (LSI) and later the very large scale integrated circuit (VLSI). This combination of discrete digital logic into a single chip came to be called an application specific integrated circuit (ASIC). It achieved the cost, size and power consumption advantages inherent in integrated circuits, but it lacked the flexibility inherent in software-driven general purpose processors.

1.1.3. Growth of the Wireless Markets and Multiple Standards

In the late 1980s and early 1990s, a variety of markets for portable two-way wireless communications systems were experiencing rapid growth and rapid evolution of technology for implementing modulation/demodulation and encoding/decoding, resulting in a proliferation of AISs. Each advancement in AISs increased the performance requirements on baseband

processing. The pace of requirement growth and the scale of the challenge is illustrated by the early development of IS-95. The first test implementations of IS-95 were completely described in code (software) because of the need to make and test quick refinements. As a result, one early prototype of an IS-95 handset consisted of 64 6800 processors rack mounted in a panel truck.

Demand for universal wireless communications began to appear, see [2]. Users in military, civil government, and commercial telecommunications began to demand 'anywhere' wireless communications and additional wireless services such as geolocation. In each market there was a large and growing number of noncompatible AIS footprints, none of which could by themselves provide the 'anywhere, anytime' service desired. Manufacturers were forced to design and build an inventory of a large number of different models and to manage the logistics of getting the right model to the right footprint at the right time to enable a customer. [6] Network operators were also facing high costs from 'fork-lift upgrades' (complete infrastructure replacements) as they moved from one generation of AIS to another.

1.1.3.1. Parallels with Computing

The situation was very similar to that in computing in the mid-1970s, where each application required a specially built dedicated 'intelligent terminal', which resulted in limited usability, deployment and management problems. The subsequent advent of the PC, with its standard hardware platform, which could be suited to a wide variety of tasks by application of the appropriate software, dramatically changed computing. By the early 1990s, wireless markets began to search for an analogous transformation of wireless equipment.

The lack of flexibility inherent in ASIC processors for high speed signal processing made equipment manufactured to support a given AIS, for the rest of its useful life, limited to that AIS. This made it difficult for commercial users to communicate as they moved in and out of different AIS service footprints. Government users who didn't have the *lingua franca* of the public switched telephone network to fall back on, if they had equipment supporting two different AISs might not be able to communicate at all. ASICs, then, satisfied the desire for lower cost, size and power consumption but didn't satisfy the desire to make devices easier to interoperate or to manage in the field.

1.1.3.2. Software-Defined Radio

In the early 1990s, solutions to the high speed signal-processing requirements that offered software driven flexibility and the ability to change AIS baseband subsystems to support different AISs through software began to appear [3, 4, 5]. The appearance of these solutions, in the company of the market drivers, led to the coining of the term software defined radio (SDR). How SDR fits into the general evolution of RF systems is displayed in Figure 1.2.

These initial solutions can be characterized as being based on reconfigurable logic. There are a variety of approaches within this general area. Generally, they use software or software-like code to configure digital logic to perform the high speed signal processing at relatively low clock rates, thereby achieving the desired power consumption/heat dissipation while they are running, while being capable of being reconfigured between runs to support different AISs [3, 4, 5].

1.2. Evolution of SDR Baseband Requirements

1.2.1. System Optimization

Over time, information processing systems have consistently sought to optimize cost/performance by managing the balance between three elements:

- local processing/local memory,
- communications capacity,
- type and quality of service.

Local processing/memory has to do with the capability to store and process information locally. It includes such factors as the ability to handle different kinds of information (voice, data, image, full motion video, etc.).

Communications capacity has to do with how much raw transmission capability is available. It may be measured in many different ways, depending upon application – number of twisted pair equivalents, Hertz, voice-grade equivalent channels, NTSC/PAL/SECAM/HDTV, number of TV channels, number of simultaneous low speed Async (Asynchronous) data channels, etc.

Type of service includes such factors as audio, image, video, Async, Isoc (Isochronous), etc. Quality of service includes such factors as error rate, availability, latency, fidelity, etc. Cost includes a broad range of factors including equipment cost, cost of ownership, migration cost, opportunity cost of consuming scarce resources, etc.

All those in the communications industry working together over time, seek to optimize systems for the highest possible types and qualities of service at the lowest possible cost. The major players in this optimization process include the equipment designers, system operators, regulators and end users.

Certain threshold phenomena come into play to define the types of service that are 'possible'. In this context, 'possible' is a function of technical feasibility with current technology, reasonable cost, within the range of what can be imagined and desired. If communications resources are considered relatively scarce (expensive) and local processing is considered available and relatively cheap, then communications systems will maximize the utilization of local processing/memory and minimize the utilization of raw communications resources. Full function PCs using low speed dial-up lines are an example of such an optimization. If, on the other hand, communications resources are considered plentiful, but processing/storage is considered scarce, then systems maximize the use of communications resources and minimize the use of local processing/memory. X Windows is an example of such an optimization.

1.2.2. Market Drivers of SDR

SDR is both a product of, and comes to market in the context of, dramatic changes in the world's information technology and usage environments.

The Twentieth Century was shaped by one-way wireless communications systems, namely broadcast radio and television. Prime examples include Franklin D. Roosevelt's radio broadcast of the Fireside Chats, Hitler's radio broadcast of rallies and the US civil rights movement reliance on TV coverage of the response to protest activity. The Twenty-first Century will be shaped by two-way interactive, multimedia, wireless communications. Early examples include the SMS-organized demonstrations that toppled a recent Philippine administration.

1.2.2.1. Usage Trends

The most dramatic consequence of this change process is the movement from the situation where a small fraction of 1% of population uses interactive wireless communications for a small fraction of 1% of the time primarily for voice services, to the situation where 90% of the population of the planet use interactive wireless communications 90% of the time for a wide range of multi media services.[†] This is making the limited quantity of radio frequency spectra a scarce resource (expensive) and is forcing an optimization shift, moving systems to apply additional local processing/memory resources which are seen as more available (less expensive). This optimization move is creating a rapidly increasing rate of technical innovation in air interface standards (AISs).

1.2.2.2. 'Anytime, Anywhere'

Not only is the quantity of usage increasing, but the users are moving to an expectation of 'anytime, anywhere' communications. Because the requirements can change so dramatically from location to location and from time to time, and since once deployed, large communications systems continue to persist in spite of the availability of new, more effective modes (for example the continued operation of the Baudot Telex System), the number of non-compatible AISs available to a particular user is large and growing.

1.2.2.3. User, Manufacturer and Operator Needs

Users don't want to carry multiple devices in order to operate in whichever AIS footprint meets their immediate (time, place, etc.) need. Manufacturers don't want to have to manufacture a wide profusion of different combinations of specific AISs and get the right combination into the right part of the world at the right time in order to make a sale. Operators don't want to have to augment one network completely with a second, and run the two networks in parallel until they can reterminalize their user base.

Hence, the drive to SDR. SDR is seen as being able to allow the user to carry a single device; the manufacturer to reduce product line complexity; and the carrier to offer integrated services across a variety of AIS infrastructures while limiting the cost of technology upgrades. All of these factors come together to increase the demands on the SDR baseband subsystem.

1.2.3. Technology Drivers of SDR

Alongside this fundamental set of market needs, also driving SDR is a series of technology changes that is also increasing the demands on the SDR baseband processor. These technology changes include the following:

- Movement from simple analog modulation, to complex modulation (from FM through QPSK to OFDM, etc.) multimedia, digital modulation, wireless systems (from 1G through 2G, to 2.5G, 2.75G, 3G, 3.XG, 4G...)

[†] In this context, usage may mean connected ready to receive and transmit in real time, not necessarily actually transmitting and receiving 90% of the time.

- Movement from a small number of large geographic area frequency reuse plans (100s of kilometers – metropolitan area radio and TV stations – to 1000s of kilometers – clear channel broadcast radio stations) – to a large number of small geographic area frequency reuse plans (metropolitan cell sites – 5 to 50 kilometers, microcellular cell sites from 50 meters to 5 kilometers; to access points 1 to 100 meters, to PAN access points from 0.1 to 2 meters).
- Introduction of MIMO (multi-in, multi-out) technology, sometimes called smart-antenna technology, which involves multiple antenna/RF/baseband paths that are combined to achieve improved spectral efficiency, signal quality, etc.
- Increased focus on overcoming climbing noise levels resulting from congestion, leading to requirements for adaptive high Q filters and other techniques either to reduce noise levels, improve signal levels or improve performance in low signal-to-noise ratio environments

When all of these factors are combined, the rate of increase in processor demands is greater than the rate of increase in processor capability deriving solely from Moore's law. This is graphically represented in Figure 1.3.

The SDR baseband subsystem must meet these rising requirement levels, seeking better end-user performance (important to the consumer), more cost effective network implementation (important to the carrier), reduced product-line complexity (important to the vendor), and improved spectral efficiency (important to the regulator, carrier, and end-user). These requirements are expressed somewhat differently depending on which of several form factor/application environments is being targeted.

1.2.4. Requirements by Application Environment

In the environments where SDR is commonly considered, applications can be broken down into two general categories – end-user and infrastructure. End-user systems generally provide




Figure 1.3 Relative increases in available and required processing power

one or a small number of simultaneous channels. Infrastructure systems generally provide many simultaneous channels supporting a large number of end-users at the same time, while also providing access to larger networks in the supported AIS, or gateways into other networks. End-user systems can be broken down into mobile and fixed systems. In the following sections we will consider the differing requirements that emerge in a variety of end-user and infrastructure applications.

1.3. Mobile End-user Systems

Mobile systems can be simple, low functionality devices or more complex systems. They can be pocket sized or relatively large vehicular-mounted systems. New categories are evolving rapidly and so any taxonomy will constantly encounter new categories. For current purposes we will consider the following categories of mobile systems:

- handsets;
- converged devices (e.g. PDA-phones);
- notebook computers and other devices of a similar size;
- vehicular systems.

These four categories can be differentiated by progressive increases in size and power consumption, going from handsets through PDAs and notebooks to vehicular systems. Handsets are designed to be as small as possible and to function without recharging the power supply for several days. They started out providing a single channel of a single AIS. What is emerging is a set of requirements that includes multiple modes and bands of cellular, WLAN (wireless local area network), WPAN (wireless personal area network) and geolocation services such as GPS (Global Positioning System). Services supported include voice, MMS (multimedia message service), data, still image (photographs) and full motion image (video). Cellular services can include first, second and third generation services (currently work is under way on 3.5 G and 4 G networks). WLAN services can include 802.11x (data, VOIP, etc.) and cordless telephony. WPAN can include Bluetooth, ZigBee, etc.

1.3.1. Simple Handsets

1.3.1.1. High Speed Signal-Processing Requirements

Single AIS handsets today require modulation/demodulation of a single RF carrier at a time. In TDMA systems, hand-off between cells is accomplished by fast switching between channels. In CDMA systems today, resources typically applied to improve signal quality (fingers of Rake filters, etc.) are reallocated to produce a system where two code division channels within a single RF carrier are handled simultaneously to assist in 'soft' hand-offs between cells.

More complex 'simple' handsets are now appearing and will soon become dominant. The number of bands supported is increasing. Triple-band handsets are now common; even five-, six-, and seven-band handsets are appearing, and the number of bands will continue to increase. Although different bands of the 'same' modulation technique appear to have the same modulation/demodulation requirements, there are differences. For example, European GSM900 (900 MHz) and GSM1800 (1800 MHz) have somewhat different requirements.

The number of modes supported in 'simple' handsets is also increasing. In the US, all 2 G standards required fall back to 1 G (AMPS). In Europe, ETSI originally required all 3 G handsets to support at least one version of GSM. In the USA, the FCC introduced a requirement for E911 (emergency mobile phone location service); in Asia and Europe the consumer location based services (LBS) are beginning to appear. The LBS requirement is being satisfied today by some vendors through the inclusion of a GPS receiver in the handset, or time of arrival (TOA) software in the handset, or other network based solutions. Support for wireless headsets and other PAN devices require yet another mode.

Some of this growing number of modes and bands can be supported by changing the functionality of hardware with software (SDR) in a sequential manner as the user moves in and out of different footprints or experiences changes in his immediate requirements. Others require simultaneous support. If the handset's software can change fast enough, some of what appear to be simultaneous channel support requirements can be achieved in a sequential fashion.

In handsets, the high speed signal-processing unit must, in all states, minimize power consumption. To this end, it must be able to reduce functionality in the standby mode to the minimum necessary to support receiving a call initiation page and sending 'keep alive' messages in the transmit mode. At the same time it must be able to switch quickly between the wide variety of modulation/demodulation schemes. It must also be able to service multiple, apparently simultaneous, channels in a potentially constantly varying mix.

The initial approach to these 'composite' requirements was with a composite processor, that is, a separate processor for each AIS. This brute-force method seems eminently workable if one relaxes the power, size and cost constraints. However, logistic problems that result from different generations of low speed signal-processing solutions can make this approach difficult. In any case, the pressures of power, cost, and size favor a single chip integrated solution. Also, the TOA analysis for a single AIS is best performed in the high speed signal-processing portion dedicated to that AIS. TOA results can be significantly improved by combining information from multiple AISs and multiple base stations.

In the past, smart antenna, or MIMO (multi-in, multi-out) technology has only been applied to base stations. Because systems requiring multiple independent RF paths and RF front-ends have been relatively large, expensive and power intensive, it has not appeared practical to implement MIMO in handsets. This situation is, however, beginning to change. There are RF front-ends in development that can provide multiple independent RF paths at low cost, size and power. Therefore, in the near future, SDR high speed signal processors will have to also support MIMO algorithm processing.

1.3.1.2. Low Speed Signal-processing Requirements

Early 2G handset low speed signal processors only had to support conversion of voice back and forth from analog to digital, with some fairly simple interleaving and forward error correction techniques added. Now, asynchronous data, music compression, and image compression must also be supported, along with location reporting and multiple simultaneous channels. Additionally, several simultaneous channels will have to appear to be handled simultaneously.

Because of the size of the displays, the resulting image resolution is likely to remain relatively low. Because of small high quality headphones, requirements for audio quality ('resolution') are relatively high. Again, because of size/power limitations, memory size will be limited thus favoring lower 'resolution' compression/decompression with concomitant

demands on low speed signal processing. The handset usage model may be structured to limit the aggregate processing demand by downloading multimedia material in compressed form and directly storing it. In this model, decoding (decompression) only occurs when the user is 'off-line', that is, not concomitant with the actual AIS processing. Coding can be handled in a similar fashion.

Asynchronous data will likely consist of short text messages and small screen adapted web pages. The small screen size makes it difficult to handle large documents on a handset, therefore, users will limit the message/e-mail usage to relatively short messages. This will produce a limited demand on low speed signal processing in terms of volume of data. Because of expectations established in the user community through use of high speed broadband wired access to the web, speed of delivery of web pages to handsets is an issue. These demands will foster requirements to speed the flow of web data through the low speed signal-processing system.

The requirements for low speed signal processing in an SDR handset, then, are to be able to code/decode audio streams (voice and music), relatively low resolution image, fast web pages and small size text messages while moving seamlessly between a variety of AISs. The challenge is in accomplishing this processing with small physical size processors consuming limited amounts of power.

1.3.1.3. Packaging

Because of the premium on size and power, there are incentives to minimize the number of separate components. Some of the incentives result from the fact that additional power is required to go off-chip (line drives, etc.) and that packaging can greatly increase the volume occupied by a single die (unpackaged semiconductor). There can also be cost implications. The result is that some vendors are packaging the entire SDR baseband solution (processor or processors, and scratch-pad memory) in a single semiconductor. Others are combining the SDR controller with the low speed signal processor and scratch-pad memory. Both of these approaches use separate nonvolatile memory stores for both user data and program storage. As the number of functions implemented in software increases, and as the number of alternative software stacks for each function increases, the size of nonvolatile storage will have to increase. How many of these alternative software stacks should be stored on the network and how many stored on the handset is a complex optimization involving latency, cost of storage, cost of communications, implications for cost/complexity of network management and business models.

Yet other vendors are focusing on packaging the controller, low speed signal processor and high speed signal processor on a single semiconductor.

As we will see later in this chapter, there are a number of alternative high speed signal-processing architecture alternatives, each with different implications for packaging. As a result, there are likely to be a variety of approaches to minimizing chip count.

1.3.1.4. RF Baseband Interface Requirements

Software efficiency (development, debugging, life-cycle management) is greatly enhanced if it is possible to establish standard APIs (application programming interfaces). This is true in general, but especially true for SDR devices. Because of the range of alternative architectures

available and the ongoing technology innovation in the SDR field, it is unlikely that the internal architectures for SDR RF front-ends and SDR baseband processors will become standardized. However, it will greatly assist network operators, equipment manufacturers and software developers if a standard interface between the RF front-end and the baseband processor can be developed. Such a standard will have to accommodate variety in number and type of channels, but should provide a standard meta language for describing and reporting a devices' exact configuration.

1.3.2. Converged Devices

Today, in addition to simple cell phone handsets, there are a wide variety of personal portable information devices including:

- PDA/organizer;
- digital camera;
- portable game console;
- walkman;
- video walkman;
- MP3 player;
- broadcast radio and TV receiver;
- satellite audio;
- navigation system.

This list is constantly growing. Today, there are converged devices that combine cell phones with different subsets of this list. For example, camera phones are popular in Japan. Nokia has launched a game/cell phone combo, and cell phones with some form of audio for music are appearing. Within a short period of time, converged devices will support all of these functions in a single device. Users will want a single device to support the ability to switch between these functions and sometimes to run these functions concurrently. This will place additional requirements on SDR baseband processors.

These devices tend to be larger than handsets, with larger batteries and with usage patterns that generate shorter times between battery recharges. These factors somewhat reduce the power conservation constraints faced by stand-alone handsets.

These devices have larger human interfaces, thus offering richer communications experiences for the user. Medium sized screens lead to medium resolution graphic capture, transmission and display. Medium sized text files and web pages become practical. One consequence of this is that, although WLAN support is not often required on simple handsets, it is becoming more common on converged devices.

Because they deliver a broad range of applications that users come to feel are critical to their everyday functioning, the demand for 'anytime, anywhere' operation is increased. This leads to a demand for support of a wider array of AISs and support for more functions simultaneously.

1.3.2.1. High Speed Signal Processing

As compared with the handset, the SDR converged device, high speed signal processor has more power available and can occupy more volume. However, offsetting this, it must process

more modes, more bands (AISs) and support a larger number of apparently simultaneous channels. This means that technology solutions that cannot meet the stringent size/power/cost constraints of the simple handset, but have somewhat more capability, may find an applications niche in the converged device area.

1.3.2.2. Low Speed Signal Processing

As compared with the handset, the SDR converged device, low speed signal processor must provide moderate resolution image and web pages while also coding/decoding a greater (possibly simultaneous) flow of information to and from local nonvolatile storage. Here again, the usage model may be structured to limit the aggregate processing demand by downloading multimedia material in compressed form and directly storing it. In this model, decoding (decompression) only occurs when the user is 'off-line', that is, not concomitant with the actual AIS processing. Coding can be handled in a similar fashion.

1.3.2.3. Packaging

As compared with the handset, the biggest difference in SDR converged device packaging is the larger local nonvolatile memory for storing both program code and application data. This requires large nonvolatile stores that cannot be practically integrated with other functions at this time.

The other major difference is the amount of application processing required. One approach is to combine the application processing and control functions on a single processor. Another approach is to attempt to use a single powerful application processor to perform signal processing as well.

1.3.2.4. RF Baseband Interface Requirements

As compared with the handset, the biggest difference in SDR converged device RF front-end-to-baseband interface is the increase in the number of modes and bands that must be supported. There is also an increase in the area available to mount antennas, and so there is likely to be a greater MIMO requirement. These factors may yield a requirement for a greater number of apparently discrete RF baseband interfaces. However, the same API meta language used for the SDR handset should be able to describe and report these interfaces thus achieving economies of scale.

1.3.3. Notebooks

Notebook computers are increasingly incorporating wireless communications as a standard feature, and some portable DVD players are also taking this form factor. This is driven by a higher bandwidth interface between the user and the device (large screen, full size keyboard, etc.). These devices also have larger batteries, shorter recharge intervals and more volume available for processors and memory. However, because of their larger size, they are not subject to as great an 'anytime, anywhere' requirement. Similarly, they are not well suited to simple voice services. These devices have larger internal nonvolatile memory stores, but they also have large removable local stores that can be used to store both SDR program data and user data.

Notebooks, are typically only required to operate for a few hours on a single battery charge and are often operated when plugged into external power sources.

1.3.3.1. High Speed Signal Processing

As compared with SDR handsets and converged devices, the SDR notebook requires fewer simultaneous channels, and a smaller palette of locally stored AISs from which to draw.

1.3.3.2. Low Speed Signal Processing

As compared with SDR handsets and converged devices, the SDR notebook requires coding/decoding of large text files, high resolution graphic data and large web pages. The user is less likely to accept the 'off-line' download usage model and is more likely to require access to graphic and audio information as it is downloaded over the AIS. The notebook user is also likely to be a source of such information for such applications as multimedia conferencing and other peer-to-peer applications.

1.3.3.3. Packaging

As compared with SDR handsets and converged devices, the SDR notebook has less of a necessity to limit component count. The high speed signal-processing requirements may be somewhat reduced, while the power/size/cost constraints are relaxed. Therefore, one approach being pursued to meet SDR processing requirements is the use of large high speed applications processors combined with 'accelerators' where necessary. Large nonvolatile memories and removable memory media are common in this form factor.

1.3.3.4. RF Baseband Interface

Because of the increased size and power available, there are SDR baseband architectures that are more practical in the notebook than in the handset or converged device. There are also differences in the noise environment in notebooks that result from the type of processors and the type of computer buses used. This may lead to RF front-end-to-baseband interfaces that are somewhat different than for the handset and converged device. However the handset meta language should be able to describe and report these interfaces.

1.3.4. Vehicular Systems

Vehicular SDR systems begin to take on some properties more common to SDR infrastructure, whilst still retaining some of the properties of individual end-user systems. Vehicular systems are evolving into integrated highly functional systems that can take advantage of the larger antennas, larger volume and larger power sources available to deliver a variety of services to a variety of users over one or more in-vehicle wired or wireless networks. In some cases, they also provide services to users outside the vehicle within a short distance.

In addition to having larger antennas, vehicular systems are allowed by regulatory agencies to use more powerful transmitters. The power of transmitters in personal devices is

limited due to health concerns about their close proximity to people. Vehicular systems still have power limitations as compared with some infrastructure systems, again because of health concerns.

Compared with all of the types of services described previously, vehicular SDR systems provide additional unique services. These systems can be grouped into transportation information systems, telemetry systems, and guidance and control systems. The boundaries between these three types of unique service are not exact and firm – one type of service may bleed into another type.

Examples of vehicular information services include wireless toll payment, traffic alerts, etc. Telemetry may be gathered and delivered to the integrated SDR hub either wirelessly or with wires. Examples of telemetry include reporting airbag deployment to remote locations or reporting a low air pressure condition of a tire on a vehicle to the driver. Examples of guidance and control systems are the use of radar-like systems to determine the distance to the next car on the highway. Because information systems and guidance/control systems have such different requirements for reliability and latency, they are sometimes implemented in physically separate systems.

1.3.4.1. High Speed Signal Processing

As compared with personal SDRs, vehicular systems require more modes and bands (AISs) and more simultaneous channels. However, there is more power and more volume available to accommodate system components.

1.3.4.2. Low Speed Signal Processing

As compared with notebooks, the low speed requirements are quite similar. The major difference is in the number and volatility of simultaneous channels.

1.3.4.3. Packaging

SDR is a particularly attractive approach for vehicular communications systems. Although the total space and power available for communications is large, if each function is implemented in a separate radio, the space and power required would quickly outstrip what is available. The ability of SDRs to support a wide range of wireless applications with a single relatively small package is extremely valuable to the vehicle designer.

The other major characteristic is that vehicular systems provide a somewhat hostile environment characterized by high temperature ranges, high G-force loading, high vibration levels and the possible presence of caustic substances. These characteristics tend to favor small ruggedized packages. This puts a premium on reducing part count and heat dissipation.

1.3.4.4. RF Baseband Interface

Because of the increased size and power available, there are SDR baseband architectures that are more practical in vehicular systems than in personal devices. There are also differences in the noise environment in vehicles that result from engines, spark plugs, etc. This may lead to RF front-end-to-baseband interfaces that are somewhat different. However, because these

services are still primarily end-user focused, the handset meta language should be able to describe and report these interfaces.

1.4. Fixed End-User Systems

Fixed end-user systems have many of the same requirement profiles as mobile systems. However, there is a relaxation of requirements in the areas of low power, small size, and multifunctionality. Fixed devices can be connected to utility power. Since they are not being carried around, size and weight is not such an issue, nor is accommodating many different types of functionality in a single package.

At the same time, the demand for mobility by itself in our society, as well as expressed through economies of scale, is driving many fixed wireless systems to become mobile. Given these forces at work, the proportion of fixed wireless systems deployed is likely to decline. That is, although there may be increases in the absolute number of fixed systems, the rate of growth of mobile systems will be much greater.

1.4.1. Infrastructure System Requirements

Infrastructure systems can consist of one or many different types of system. For example, large cellular network infrastructures consist of basestations which are connected via an internal point-to-point network link to a basestation controller, which is in turn connected to other controllers, such as separate addressing, billing and maintenance systems. On the other hand, simple WLANs often consist of an access point in a computer connected to a wired network. In this section, we will be primarily concerned only with that portion of the infrastructure that services personal or vehicular devices.

There are two major factors that differentiate SDR infrastructure systems from mobile end-user systems. The first involves SDR usage models and the second involves inherent functionality.

Mobile end-user systems are constantly moving in and out of noncompatible AIS footprints. This requires these types of systems quickly to change from one to another AIS. They also support simultaneous concurrent noncompatible AISs, so these types of system may be required to maintain some AIS channels while reconfiguring others. In this respect, mobile personal systems can be characterized as having a relatively high volatility of AIS change/reconfiguration.

Infrastructure systems are stationary by nature. The predominant usage model is that mobile devices will change to configure to the AIS supported by the relevant infrastructure as they move in and out of infrastructure footprints. Therefore, except in unique specialized cases, infrastructure systems are not supporting a variety of constantly changing AISs. Rather, the SDR usage model focuses on graceful technology migration from one AIS to another. These changes, rather than being measured in fractions of a second, are measured in days, months or years. However, as new AISs are developed and deployed, change requests (network operators requests for features and capabilities different from what was described in the original purchase contract) in the order of 100 per month per infrastructure product are appearing.

The inherent functionality difference is that infrastructure systems are required to support multiple mobile devices. Therefore, infrastructure devices typically process blocks of channels rather than individual channels.

1.4.2. Basestation Requirements

Radio systems have sought from the earliest implementations to make efficient use of the RF spectrum. This has been accomplished by frequency/geographic segmentation to allow frequency reuse. As RF systems have evolved from predominantly point to point to broadcast to interactive, the geographic space assigned to given frequencies has contracted. With the appearance of cellular systems, the concept of the cellular basestation was defined. Initial basestations were implemented as concrete buildings containing rack mounted equipment with utility provided power, adjacent to dedicated antenna towers. Each basestation serviced a relatively large area. For example, typical macro-basestation cell sites ranged from 50 to 100 kilometers.

As basestations evolved, they began to be packaged in smaller form factors and serve smaller cells – microcells. Today there are still macro basestations. In addition, however, there are smaller packages that can be mounted in a utility room on top of a building. Finally, there are small weatherproof containers with heat dissipation capacity that can be mounted on a telephone pole with attached antennas. These are commonly referred to as ‘shoe box basestations’.

During this same period, WLAN access points began to appear. As such access points proliferated from being a relatively rare, end-user-deployed device, to provide carpets of footprints covering campuses, and filling hotspots for carriers, they began to take on the characteristics of ‘picocellular’ basestations.

One of the key aspects relative to SDR basestations is the extremely high requirement for reliable fast changeover. In the case of a single mobile device, a failure in a software download and install that requires a reload is not catastrophic. In the case of a network consisting of 1000 basestations serving 100,000 users, such a failure could be considered extremely serious. There have been examples of networks (both in the case of wired Internet access and wireless networks) where software changes (not associated with SDR) have caused major disruptions. In these cases, the network operators suffered significant immediate financial losses and longer term damage to their reputations for quality and reliability. When an SDR basestation is migrated to a new or modified AIS, the interruption in service if any, must be minimized.

1.4.2.1. High Speed Signal Processing

Requirements for SDR basestation high speed signal processing are more demanding than mobile systems in that basestations must modulate/demodulate relatively large blocks of channels. At the same time, as compared with the handset, the SDR basestation high speed signal processor has more power available and can occupy more volume. The processing load can vary dramatically from one AIS to another.

The introduction of smart antenna technology is creating another vector requiring SDR support. For example, one usage model is as follows: A network operator installs a base station in a suburban area that has recently been connected to high capacity transportation systems. Population and subscriber penetration is very low, but is expected to grow. As the area becomes more intensively used, population density and subscriber penetration grows. The operator wishes to increase the traffic handling capability of the basestation by adding smart antenna capability. Again, this is a graceful migration requirement.

These requirements combine with equipment and logistics costs. Providing more processing power than is needed for a particular AIS/smart antenna combination in order to provide a future migration path can have significant cost implications. On the other hand, savings associated with not having a wide variety of physical configurations deployed in a single network operator are also significant.

1.4.2.2. Low Speed Signal Processing

Here again the major differences in requirements for SDR basestation low speed signal processing and mobile system processing is that basestations must decode/code relatively large blocks of channels. These channels interface with a variety of formats including circuit switched voice, SS7, TCP/IP, etc.

In many cases, the basestation can avoid having to decode fully, individual streams of end-user data. Instead, the basestation decodes only the coding (if any) added to overcome unique wireless noise environments, and merely passes the rest of the data to another network component until it arrives at the destination end-user where it is finally fully decoded.

Over time, the back-end network to which the basestation hands off traffic may change, thus changing the coding/decoding requirements. Therefore, SDR requirements include the ability of basestations to migrate gracefully between different formats and mixes of formats as the network operator migrates their network back-end structure.

1.4.2.3. Packaging

Larger SDR basestation baseband processors are generally implemented on printed-circuit boards that can be plugged into backplanes. Smaller shoe-box basestation baseband processors are migrating to 'single board' packaging. This process is analogous to the movement from mini-computer backplane architectures to PC single-board architectures.

1.4.2.4. RF Baseband Interface

Different AISs may lend themselves to different architectures for high speed signal processing. For example, 2G basestations typically interface to the RF front-end through a single transmitter and a single receiver channel that delivers the complete band being serviced to the high speed signal processor. As WCDMA moves to 5 MHz (and beyond) baseband bandwidth blocks, some are considering separate interfaces to each block. This allows the processing load to be parsed across three, dedicated, high speed signal-processing subsystems. Similar approaches are being considered for multiple RF interfaces supporting multiple smart antenna RF baseband channels. This can lead to a requirement for reconfigurable RF/baseband interfaces to support migration to different interface architectures as AIS migration occurs.

1.5. Today's SDR Baseband Technology Approaches

1.5.1. Technologies and Criteria

There are (at least) nine technology conceptual approaches to implementing SDR baseband subsystems that can be seen today. These are:

- general purpose processors with specialized software;
- high speed DSPs;
- multiple ASICs;
- parameterized hardware;
- switchable microcode;
- multiprocessor array (including VLIW);
- reconfigurable logic;
- analog signal processing;
- combinations of the above.

Each of these technology approaches has strengths and weaknesses. Each can be characterized by a profile of the following parameters [7]:

- power consumption;
- size;
- cost;
- field upgradeability;
- silicon evolution;
- tools maturity;
- technology risk;
- time to market (TTM);
- staff availability.

Each of the technology alternatives can be characterized as having strengths and weaknesses in each of these different areas. Breaking down the analysis into these separate areas can be helpful in determining the suitability of a particular technology to a particular application. Such an analysis has been performed [7]. In the following, the most significant aspects of the alternatives being given the most attention are examined in an effort to predict the likely direction of a generalized solution.

The most challenging portion of the baseband processor continues to be the high speed signal-processing subsystem. Although the requirements in the low speed signal-processing portion have become more demanding through the addition of simultaneous channels and new modes, it is the high speed modulation and demodulation functions that remain the most demanding.

When we look at SDR baseband technology alternatives in the context of the spectrum of requirements sets, the fundamental challenge is power, size, heat dissipation and cost. Because of the way that the requirements sets are evolving, as the size/power/heat/cost constraints are relaxed, the processing demands increase.

1.5.2. Capabilities and Constraints

1.5.2.1. Conventional DSPs

The oldest, most mature, easiest to staff and shortest time-to-market solution is to use a monolithic single stream instruction set DSP. The problem is that in order to meet the system requirements, the clock speed requirements are either unattainable or produce results which don't fit the size/power/heat/cost constraints.

Power consumption in CMOS (the lowest cost/power consuming semiconductor process generally available) is directly proportional to the combination of the number of transistors and the speed with which they move through different states (clock speed). Heat dissipation is a function of power consumption. Cost is a function of die size (the size of the raw semiconductor when it is cut out of the wafer on which it was produced and before it is packaged).

To illustrate the situation, let's go back to the handset example from the beginning of the chapter of 100 instructions per sample (number of instructions per sample are increasing due to the evolution of AISs) for the modem subsystem resulting in a clock speed in excess of 1 GHz. Performing these high speed signal-processing functions in a DSP gives us a solution that has many advantages, including being fully software defined and thus fully software flexible, but doesn't meet the power/heat/cost constraints. Now add the increasing processing demands and the problem becomes more severe.

1.5.2.2. Evolved Submicron DSPs

It has been postulated that, as semiconductor technology advances and feature sizes and resulting voltages are reduced on chips, DSPs will be able to meet the constraints. However, limitations in voltages, bias currents, etc., are creating an asymptote in power consumption (and related heat dissipation). The result is that this architecture can't meet the constraints. This defines one limit case.

1.5.2.3. Optimized Hardwired Logic

Another end of the spectrum is to implement each of the 100 modem instructions in a separate piece of fully optimized hard wired logic, then string them together in the order of processing (ASIC). The theoretical result is a modestly sized die running at 4 MHz. This architecture is characterized as highly scalar and deeply pipelined. It clearly meets the size/power/heat/cost constraints. However, it doesn't meet the SDR flexibility requirement. This defines another boundary case.

1.5.2.4. Software Configured FPGAs

The first breakthrough in SDR came with the use of FPGAs. Under software control, FPGAs could be configured in a similar highly scalar deeply pipelined fashion. Unfortunately, in order to achieve software reconfigurable flexibility, FPGAs require more transistors to implement a given logic function than does hardwired logic. There are also challenges in stringing the 'instructions' together in the correct order. These problems are characterized as place and route difficulties. This architecture does achieve the desired software programmability; however, it pays for it in bigger dies that cost more and burn more power. Some maintain that the power differential can be as much as 50 times that of ASICs. Early SDR commercial basestations and military devices were implemented this way because they could relax the size/power/heat/cost constraints. FPGA technology is not as mature as DSP technology so there are also costs in development parameters (cost, TTM, risk, etc.). This defines another limit case.

1.5.2.5. 'On-the-Fly' FPGA Reconfiguration

Some maintain that innovative FPGA routing structures plus reconfiguration times can be decreased to the point where one part of the chip can be reconfigured fast enough for only a limited number of 'hard wired' logic instruction capable hardware elements to be needed. In this architecture, while one element is processing, the next element is being configured for the next instruction, and then the one after is ping-ponged back to the first element. If this were done, the silicon area would be reduced but the clock speed would be increased. The final result in size/power/heat/cost for such an architecture is not clear because, to date, none has been publicly demonstrated.

1.5.2.6. Multiple Processor Arrays

Another approach is to have multiple instruction set processors. At the limit, there would be one for each of the 100 instructions strung together. Now, the resulting system could lower the clock speed 100-fold, but if the processors were the same size as the single stream instruction set processor, the die size would increase 100-fold (or more because of the routing problem). This is another limit case.

1.5.2.7. Optimized Processors

Another approach is to optimize the processors in this array. One way to do this is to use reconfigurable microcode. In this architecture, special machine instructions (CISC on steroids) are created that are optimized for functions required to support a particular AIS. Now time/power/space resources must be used to store, determine the correct instantiation of, retrieve, and load the correct microcode at the correct time. This also consumes silicon space, time power, etc. (instruction set resource management). One way of reducing the resource management problem is to have all processors load an instruction (or microcode) at the same time. This is often referred to as a single instruction multiple data (SIMD) architecture. If the algorithm being implemented lends itself to this kind of processing (for example a series of FIR, finite impulse response, filters that are the same and connected in series), it may increase the efficiency of the resulting processor. This is only practical if the AIS can be represented by an algorithm which fits the constraints. If so, in this way, it may be possible to reduce the number of instructions tenfold and thereby reduce the clock speed. This is another limit case.

1.5.2.8. Parameterized ASICs

Finally, it is possible to create ASIC blocks that are tailored to one very particular function, but within that function to allow for software control of a particular parameter. For example, a FIR filter that can be programmed for one or another limited number of different taps, or a rake filter which can be programmed for a limited number of different forks.

1.5.3. A Way Forward?

For several years now, proponents of each of these different architectures have been competing in the SDR space. Each has sought to optimize their own particular solution to the




Figure 1.4 SDR baseband processor solution space

point where it can meet the full set of SDR requirements. Many observers now believe that the truly successful solution path is not the optimization of one particular limit case, but rather some combination of the architectures that lies within the interior of the space defined by the limit cases. The relationship of the limit cases is graphically represented in Figure 1.4.

For example, a limited number of small single stream instruction set processors each with reconfigurable microcode (multiple instructions – multiple data architecture), connected with a small set of parameterized ASIC blocks by a reconfigurable interconnect matrix, might produce the optimal processor architecture. Such an SDR baseband architecture might have several different instantiations, one optimized for each of the different AIS/application requirement sets. This optimal solution space is represented in Figure 1.4.

One of the key challenges in creating innovative hardware architectures like these lies in the programming models employed and the resulting availability of tool sets and qualified programmers. The best imaginable solution would be one that could make full use of existing (without requiring any new ones) programming models and tool sets. If this is not fully achievable, the closer it can be approached the better.

1.6. Conclusions

In this chapter we have traced the evolution of SDR baseband processing, discussed the range of requirements for SDR baseband processing systems in different form factors and applications, and considered some of the technologies that are available to implement SDR baseband processing subsystems. We have seen that evolution of AISs and usage models is creating processing demands that are increasing faster than Moore's law increases in semiconductor performance. This has led us to the conclusion that SDR solutions require innovations in processor architecture. Although some are strong proponents of a single architectural approach, given the evolving requirements and the technologies available to meet them, it is possible that the optimal solution will be found in a combination of available processing architectures.

Acknowledgments

The author would like to acknowledge assistance in building the author's understanding of the subject matter from Professor Emeritus Michel J. Flynn, Professor Tadao Nakamura, Professor Martin Morph, Ravi Subramanian and John Watson.

References

- [1] M. Cummings and S. Heath, 'Mode Switching and Software Download for Software Defined Radio – The MMITS Forum Approach', *IEEE Communications*, February, 1999.
- [2] *Software Defined Radio Forum Technical Report (TR) 2.1*, Edited by SDR Forum Technical Committee (Chair M. Cummings), November 1999, SDR Forum, NY, USA.
- [3] M. Cummings *et al.*, Method and Apparatus For Communicating Information, USA Patent No. 5 907 580, 25 May, 1999.
- [4] M. Cummings *et al.*, Method and Apparatus For Communicating Information, USA Patent No. 6 175 589 B1, 16 January, 2001.
- [5] M. Cummings and S. Haruyama, 'FPGA in the Software Radio', *IEEE Communications*, March, 1999.
- [6] M. Cummings, 'A Multimode SDR radio frequency implementation', Chapter 3 in *Software Defined Radio: Enabling Technologies*, W.H.W. Tuttlebee (Ed.), John Wiley & Sons, Ltd, Chichester, 2002.
- [7] 'Puzzle of software defined radio has multiple solutions', *Electronic Engineering Times*, 16 October, 2000.

