全国优秀出版社

科技英语教程

马新英 孙学涛 林 易 编著

国防工业出版社
National Defense Industry Press

H31
1166

科技英语教程

马新英　孙学涛　林易　编著

国防工业出版社

·北京·

内 容 简 介

本书是根据《大学英语教学大纲》中关于“理工科院校的英语教学应在两年的基础训练之后，在第五至第七学期开设必修的专业阅读课”的规定，并结合理工科院校的主要专业而编写的一本科技英语教材。全书共分三大部分，各部分由若干单元构成，分别对计算机、电信、数学等方面的知识进行系统介绍。各单元包括课文、生词表、综合练习和阅读材料四部分。

本书可作为计算机、通信、数学等专业的教材，也可作为相关工程技术人员的参考用书。

图书在版编目(CIP)数据

科技英语教程 / 马新英等编著. —北京：国防工业出版社，2005.8

ISBN 7-118-03970-5

Ⅰ.科… Ⅱ.马… Ⅲ.科学技术-英语-教材 Ⅳ.H31

中国版本图书馆 CIP 数据核字(2005)第 059637 号

国防工业出版社出版发行

（北京市海淀区紫竹院南路 23 号）

（邮政编码 100044）

北京奥鑫印刷厂印刷

新华书店经售

*

开本 710×960 1/16 印张 18¾ 338 千字

2005 年 8 月第 1 版 2005 年 8 月北京第 1 次印刷

印数：1—5000 册 定价：28.00 元

（本书如有印装错误，我社负责调换）

国防书店：(010)68428422 发行邮购：(010)68414474

发行传真：(010)68411535 发行业务：(010)68472764

前　言

根据高等学校理工科本科《大学英语教学大纲》的规定：理工科大学的英语教学要在两年的基础训练之后，进行必修的专业阅读训练，以保证英语教学的延续性，并培养学生以英语为工具阅读有关专业书刊，获取专业信息的能力。本书就是依据此规定编写的科技英语教材。本书可作为计算机、通信、数学等专业的教材，也可作为相关工程技术人员的参考用书。

本教材的编写主要着眼于以下三点：首先，选编专业内容尽可能体现出从基础理论到应用的专业知识结构，以扩大学生的专业词汇量；第二，在阅读专业资料的同时，注重阅读、翻译、写作等技能的训练；第三，从科技英语文体形式的角度选编内容，以便为学生今后进一步扩大阅读范围打下基础。

以上述三点为指导，全书共分三大部分，各部分由若干单元构成，分别对计算机、电信、数学等方面的知识进行系统介绍。各单元包括课文、生词表、综合练习和阅读材料四部分。综合练习主要包括翻译、写作、构词法等方面的内容。阅读部分在专业知识方面与课文紧密相联，每篇阅读材料后都附有生词表。本书的材料均选自英文原版书刊，涉及前言、绪论、论文、说明书等多种文体。

本书的第一部分由孙学涛编写，第二部分由马新英编写，第三部分由林易编写。全书由马新英统稿。

因编写时间仓促，编者水平有限，对于书中所出现的缺点、错误，恳请读者批评指正。

编　者

2005 年 5 月

Contents

此为试读,需要完整PDF请访问: www.ertongbook.com

Part I : Computing

Unit One

Text

Making Sense of 64-Bit Processors

In 2000, Intel introduced the Itanium processor, a 64-bit Explicitly Parallel Instruction Computing (EPIC) processor initially intended for the same types of functions as its competition. According to Intel, the Itanium processor offered up to twice the transaction performance of RISC processors, at about half the price per transaction. As Jerry Braun, product line manager for Intel's dual Xeon processors, puts it: "The Itanium replaced aging RISC." Compaq, Dell, and IBM are among the vendors currently offering servers based on the Itanium.

However, the first Itanium chips suffered from some major drawbacks. They weren't designed to run 32-bit applications, and Itanium-based products attracted a disappointingly small number of buyers. This led to the perception that Intel did the 64-bit processor "wrong".

By contrast, AMD did things "right" with its Opteron processor, released in April 2003. The Opteron was designed to run 64-bit applications, as well as provide all the benefits of 64-bit processing. But AMD went a step further. It created 64-bit extensions in the Opteron's core that allowed the chip to run 32-bit applications under a 32-bit OS. And since the Opteron could also run 32-bit applications natively under a 64-bit OS, this made the processor fully compatible with 32-bit applications. In effect, the Opteron was a 64-bit processor able to do 32-bit computing.

Intel responded by adding a 32-bit emulator to its subsequent Itanium chip, the Itanium 2. This emulator didn't run 32-bit x86 applications as native applications. Instead, it translated the 32-bit instructions into code that could be run on the Itanium chip. This extra step meant that applications didn't typically run as quickly

as applications with native support. (However, the higher performance of the newest Itanium 2 processor today can arguably compensate for the performance hit taken by translation.)

AMD followed the Opteron with a mobile version of the 64-bit processor. The company now also offers versions designed for low-end servers and workstations under the Athlon 64 family. The ability to run 32-bit applications natively (and the inability of Itanium to do the same) has helped AMD's Opteron sales to quickly eclipse those of the Itanium.

In July, Intel announced its answer to the Opteron. Previously code-named Nokona, the Xeon EM64T processor runs 32-bit applications natively and provides the ability to run in 64-bit mode. The Xeon EM64T also implements 64-bit extensions compatible with those developed by AMD.

The move to 64-bit processors is an evolutionary step that's been on product roadmaps for both Intel and AMD for some time. In the future, the chipmakers plan to phase out 32-bit processors and eventually only offer 64-bit processors. AMD, for example, is looking to produce 32-bit processors at the end of next year, but only if the market still demands them.

Why 64-Bit?

One of the primary advantages of 64-bit processors is their ability to effectively remove the memory addressing limit inherent in 32-bit processors. A 32-bit processor can only address 2^{32} bytes, or 4Gbytes, of data. While this may be adequate for basic computing operations, it's often nowhere near adequate for processing extremely large databases, complex drawings, video, multiple applications, or complex gaming.

On the other hand, 64-bit processors can access 2^{64} bytes, or 16 Exabytes, of data. This astronomical number has been estimated to be five times as large as the number of words spoken by all people throughout history. The limit on addressable memory space goes away when a 64-bit processor is running in 64-bit mode.

Intel's Itanium and Xeon EM64T and AMD's 64-bit processors are full 64-bit implementations, with 64-bit physical addressing, according to Nick Carr, marketing manager at Red Hat. When run using 64-bit OSs, 64-bit applications don't have to deal with the issues and limitations of addressable space that 32-bit applications must deal with.

64-bit processors also have additional registers, reducing the need for applications to go to slower main memory. These registers retain data being processed and deliver it to the CPU when needed. This saves the CPU from having to fetch it from cache, slower main memory or, slowest of all, hard disk.

In addition, a portion of the 4Gbyte address space provided by 32-bit processors is devoted to managing the processes being run. In Windows, this is half of the virtual address space, or the upper half of the 4Gbytes.

The benefits of 64-bit processors extend beyond the nearly unlimited memory addressing space. The processors can also fetch 64 bits of data at a time. All three 64-bit families—the Itanium, the Xeon EM64T, and AMD's 64-bit processors—support a wider data path, handling twice the amount of data that can be handled at a time by their 32-bit counterparts.

The combination is greater than merely the sum of the parts. With their ability to handle more data and to address a virtually unlimited address space, 64-bit processors can process considerably more data and run more applications than a 32-bit processor or OS.

Other functions built into the chip design also enhance performance. This is especially true for the Itanium and Itanium 2 processors, which are designed to run a limited set of instructions and to run them in parallel. Parallelism is integrated into the Itanium's architecture and allows it to fetch up to three instructions at once. The processor can execute up to 20 instructions in parallel and includes 128 registers (32 static and 96 rotating). The result is extremely fast performance for certain types of applications that lend themselves to parallelism, such as large databases, computer-aided design, ERP, and business simulations. According to a document obtained from Intel's labs, the Itanium also has the ability to provide the "illusion of infinite physical registers," removing register-related processing restrictions.

Finally, the Itanium can easily outperform the Xeon EM64T and the AMD 64-bit processors when running applications that make use of the its integrated hardware features.

Software Support and Conversion

For years, 64-bit OSs have been available for 64-bit RISC processors. These OSs included Solaris and HP-UX, among others. 64-bit OSs have also been available for the Itanium processor for some time now.

A 64-bit Linux kernel has also been developed for the AMD 64-bit processors, the Itanium, and the Xeon EM64T. This kernel has been rolled into Red Hat Linux, Fedora Linux, SUSE, and other Linux distributions.

"If you put 64-bit Red Hat OS on a box, it will run 32-bit applications and 64-bit applications. The AMD 64 and EM64T will run an old 32-bit OS or a new 64-bit OS," notes Red Hat's Carr. However, "If you put a 32-bit OS on it, you essentially throw away 32-bits of silicon—there's everything 32-bit about. If you use a 64-bit OS, then you use all the features. You can choose to run 64-bit applications or 32-bit applications."

Microsoft's delivery of a 64-bit versions of Windows XP Professional and Windows 2003 server has been pushed back to the middle of next year. However, the company has made beta versions of its 64-bit OSs free for download on its Web site. In late August, it also announced a free program to exchange 32-bit versions of Windows 2003 Server and XP Professional for 64-bit versions of the programs when they become available.

If that's too long to wait, Linux and its variants are already available. As Sun's Wambaugh notes, "The Opteron runs Linux screamingly fast. It runs Solaris screamingly fast." The same can probably be said for Linux and Solaris on the Xeon EM64T because the underlying instruction sets are extremely similar.

If a 32-bit application won't benefit from the added address space available to 64-bit applications running on a 64-bit OS, porting the application to a 64-bit version may not be necessary. These applications should run well in the 32-bit space provided by the 64-bit OS. In particular, 32-bit applications running on the AMD 64 or Xeon EM64T processors under a 64-bit OS will run faster than if they were running on a 32-bit OS. Also, more 32-bit programs can run under a 64-bit OS than a 32-bit OS.

In order to run a 32-bit application as a 64-bit application, it must be recompiled for a 64-bit OS. AMD, Intel, and third-party compiler companies offer tools that simplify the recompilation of applications written in a number of programming languages. In many cases, the recompile can be done relatively quickly, with minimal recoding.

Overall, developers aren't expected to have significant challenges recompiling 32-bit applications to run as 64-bit applications. However, there may not be much pressure to make these changes to mainstream applications because the Xeon

此为试读,需要完整PDF请访问: www.ertongbook.com

EM64T and AMD 64-bit chips handle 32-bit applications so elegantly already.

The Itanium is a somewhat different story. Because it was developed for extremely high-performance processing of very large data sets and has a unique set of registers, instructions, and other elements, compiling a 32-bit application to run on the Itanium may pose more challenges. Itanium compilers should make good use of all the Itanium's extra registers and other performance-enabling capabilities.

Many applications have been developed in 64-bit, in part because the Itanium can typically run applications that have no original 32-bit counterpart, and because the processor wasn't originally able to run 32-bit programs.

Mark Brownstein

10/01/2004 12:00 H EST

http://www.networkmagazine.com/shared/article/showArticle.jhtml?articleId=47205169

Words and Expressions

astronomical 天文学的
compatible 兼容的
compiler 编译器
eclipse 超越，使黯然失色
emulator 仿真器
exabyte 百万万亿字节
hard disk 硬盘
implementation 实现
instruction 指令
main memory 主存储器
parallelism 并行性
perception 观念，概念
programming language 程序设计语言
recode 重新编码
recompile 重新编译
register 寄存器
transaction 处理，事务
vendor 卖主

Exercises

Reading comprehension

I. Answer the following questions according to the text.

1. What are the disadvantages of Intel's Itanium processor?

2. AMD's Opteron is considered better than Intel's Itanium. What's the reason?

3. What is the function of the emulator in the Intel Itanium 2 processor?

4. Why is it necessary to remove the memory addressing limit inherent in 32-bit processors?

5. What's the significance of integrating parallelism into the processor architecture?

6. What are the measures taken by Microsoft to support 64-bit processors?

7. Should all mainstream applications be recompiled to run as 64-bit programs?

8. Will you equip your PC with a 64-bit processor?

Word-building

II. Summarize the following prefixes.

monoflop	octal
unidirectional	decimal
ambipolar	hexadecimal
bistable	microvolt
duplex	nanosecond
dipolar	kilobyte
quadratic	megabyte
tristate	gigabyte
quadrant	terabyte
Pentium	exabyte

Translation

III. Translate the following sentences into Chinese.

1. A data block is the unit of transfer between main memory and auxiliary storage and usually consists of several records.

2. We define a hacker as a "computer enthusiast who enjoys learning everything about a computer system by means of clever programming".

3. Data structures which change in size once they have been created are called dynamic data structures.

4. As the use of computer networks, especially the Internet, has become pervasive, the concept of computer security has expanded to denote issues pertaining to the networked use of computers and their resources.

5. Multiprogramming means the existence of many programs in different parts of main memory at the same time.

6. In general, integrity refers to maintaining the correctness and consistency of the data.

7. Task management routines, which operate in the supervisor state and are primarily core-resident, are frequently referred to as the supervisor of the operating system.

8. A facility allowing a peripheral device to transfer characters directly to

memory without going through the CPU is known as direct memory access (DMA).

IV. Translate the following sentences into English.

1. 目前，IBM、戴尔等厂商都在推销基于 Itanium 处理器的服务器产品。

2. AMD 发布的 64 位处理器 Opteron 是针对 64 位应用的，但也完全兼容 32 位应用。

3. Intel 的 Itanium 处理器是全 64 位的，不再有 32 位处理器固有的寻址空间限制。

4. Intel 的 Itanium 处理器一次能取 64 位数据，最多可并行执行 20 条指令。

5. 更多的寄存器也使 CPU 省却了从高速缓存、内存或硬盘取数据的麻烦。

6. 在 32 位处理器的 4GB 地址空间中，有一部分是专用于管理正在运行的进程的。

7. 微软公司已把其 64 位操作系统的 β 测试版放到了网站上，可免费下载。

8. AMD、Intel 和一些第三方公司现在都提供一些工具，用以简化多种语言应用程序的重编译工作。

Reading

Platform 2015 Unveiled at IDF Spring 2005

"Imagine a phone that can translate languages in real time..., or finding a photo of your children playing with a pet from among thousands of photos... To deliver these capabilities in products that are easy to use and attractive to many people requires that we, as an industry, rethink our approach to platform development."

—**Justin R. Rattner**, Intel senior fellow and director of the Corporate Technology Group

As the new director and public face to Intel's research and development endeavors, Justin Rattner articulated Intel's vision for a more natural and humanized form of digital intelligence in his keynote at the Intel Developer Conference (IDF) Spring 2005 conference in California.

In the Platform 2015 vision, technology is not only more natural and easier to use, but also less prescriptive and more predictive. Reiterating concepts from the Era of Tera keynote first put forward by Pat Gelsinger at IDF Spring 2004, Rattner described a time when technology will rapidly and transparently synthesize high volumes of complex data and explore different outcomes, ultimately helping people

do what they want to do more easily.

Rattner illustrated some of the potential future usage models that Intel is experimenting with to understand their impact on future architecture.

Technology that People Want to Use

Rattner outlined how researchers at Intel are studying the types of tasks people will want their electronic products to do. The researchers then use that knowledge to drive hardware and software technology development that will serve as the foundation for more intelligent future platforms.

"We should have interfaces that understand that, as a human, we're fairly error-prone," said Rattner. "You know, 'to err is human.' ...Instead of getting a nasty 'File not found' message, we'd like to hear, 'Did you mean this? Did you mean that?' Much like Google does today, but expanded hundreds of times." He continued,"Imagine... that your home actually helps you live a healthier lifestylethat it monitors (your health)...and detects signs of an early disease, ...provides information to your physician, and generally gets those problems corrected early on. Imagine a phone that can translate languages in real time so you can talk to people in other countries more easily, or finding a photo of your children playing with a pet from among the thousands of photos you have stored in multiple computers in your house."

Even Hollywood joined Rattner to support Intel's Platform 2015 vision, in the form of Steve Sullivan, director of research and development at Industrial Light and Magic (ILM). ILM is the digital effects company behind the *Star Wars* and *Pirates of the Caribbean* movies. After wowing the crowd with samples of digital effects in virtual-reality pieces, Sullivan repeated the need for greater amounts of processing power to meet the ever-increasing demand for exciting and realistic digital content, whether for entertainment or electronic commerce.

In order to create the technology and platforms that support many of these usage models, processing capability needs to increase by hundreds or even thousands of times. "(Some of these) tasks might seem simple," Rattner explained, "but they require levels of performance, sophistication and intelligence in both hardware and software that don't exist today. To deliver these capabilities in products that are easy to use and attractive to many people requires that we, as an industry, rethink our approach to platform development."

New Uses Require Architectural Innovation and Evolution

In order to support new usage models, future processors and platforms based on them must evolve through innovations in many key areas, including:

Low power

Reconfigurable architecture

Accelerators

Advanced wireless communications

System management

Security

Identity

Rattner focused on two technology areas fundamental to the Platform 2015 vision: virtualization and large-scale, thread-level parallelism. 3D-intensive applications, such as games, have posed a unique challenge to virtualization because they dominate the graphics pipeline and prevent multiple virtual machines from sharing the physical graphics device. Another R&D engineer joined Rattner to demonstrate a prototype solution that allows an operating system and graphics driver to run in a virtual machine in the same way that it does on a "bare metal" machine. This capability allows for different virtual machines to share the physical hardware. This effective sharing can enable multiplayer gaming on one machine and enable other display-intensive applications to run within multiple windows.

Rattner introduced the "many-core" concept, explaining that Intel researchers and scientists are experimenting with "many tens of cores, potentially even hundreds of cores per die, per single processor die. And those cores will be supporting tens, hundreds, maybe even thousands of simultaneous execution threads."

Software, particularly large-scale, thread-level parallel processing, is the first challenge in making the many-core platform a reality. In addition to developing new parallel programming tools and libraries, Intel is working on domain-specific parallel programming. In conjunction with the Chinese Academy of Sciences and the University of Texas at Austin, researchers are striving to establish a programming environment for packet processing that yields near-equivalent performance to hand-tuned code on highly threaded architectures using Baker, a domain-specific programming language. Once successful with a few domains, research will turn to more general-purpose parallel programming.

Key Challenges and Emerging Solutions

Stacked Wafers and Stacked Dies—Progress in Improving Memory Bandwidth Rattner noted that while addressing one problem, others can get introduced. Such is the case with the need to scale memory bandwidth to keep up with the increasing computing capabilities possible with multi-core and many-core platforms. He illustrated stacked wafers using "thru-silicon vias" and stacked dies—two novel packaging techniques that are showing promise in meeting the memory bandwidth challenge.

Usually, vias are the holes in printed circuit boards used to electrically connect different layers. Intel is experimenting with applying the same principle to silicon chips—to stack wafers on top of each other. Wafer stacking can be mass-produced since all of the dies on the wafers are aligned and bound with one step. However, since the dies are cut after the wafers are fully stacked, each die must be the same size—and that can introduce inefficiency. Die stacking has the advantage that dies of different areas, pitches, and technologies can be combined. It is also possible to combine logic with analog components. However, it is not as readily mass-produced as wafer stacking.

While Intel research and advanced development teams continue to work on this challenge, Rattner noted,"if we can get a four or five orders of magnitude increase in connectivity, we've made a tremendous step forward in addressing the memory bandwidth challenge."

Silicon Photonics and Chip-to-Chip Signaling Another important challenge is improving chip-to-chip signaling, as well as signaling between other elements of the platform such as storage, network, display, and other I/O devices. Intel is optimistic about bringing optical signaling onto the platform, particularly given the latest breakthrough from Intel's silicon photonics research team—the world's first continuous silicon laser.

Beyond Intel's desire for higher bandwidth chip-to-chip connections, silicon photonics can be used for network and display interconnects on desktop, mobile and entertainment platforms, backplanes in server racks, and data center fabrics. Going beyond these more traditional uses, it may be possible to use silicon lasers in surgery and for chemical analysis with the goal of having a "lab on a chip." "There are vast new opportunities here with (this) breakthrough technology," Rattner proudly stated.

此为试读,需要完整PDF请访问: www.ertongbook.com