The coordination of investment in systems of complementary assets: a clinical study of inter-firm and intra-firm mechanisms[◊]

Peter Miller[a*], Ted O'Leary[b]

[a]*London School of Economics and Political Science*
[b]*University of Manchester, UK, and University of Michigan*

Not for citation or further circulation

[◊]This study was made possible by the support and cooperation of many employees of Intel Corporation. We are indebted particularly to Gerry Parker and Mike Splinter for granting permission for the study, and to David Layzell for his support and encouragement throughout. We are grateful for comments received from participants in research colloquia at INSEAD, the London School of Economics and Political Science, Michigan State University, University College Dublin, the University of Manchester, the University of Southern California and the University of Ulster at Jordanstown. We are grateful also for comments received from Shannon Anderson, Ian Garrett, Sue Haka, Bob Kaplan, Joan Luft, Mike Newman, Brian Pentland, John Roberts, Geoff Sprinkle, Norman Strong and Yanling Zhang. The financial support of the British Design Council, the Centre for Analysis of Risk and Regulation, and PriceWaterhouseCoopers is gratefully acknowledged.

* Corresponding author. Department of Accounting and Finance, London School of Economics and Political Science, Houghton Street, London, WC2A 2AE, UK. Tel: +011 44 207 955 7231. *Email address:* p.b.miller@lse.ac.uk

Draft:November 25, 2002.

**The coordination of investment in systems of complementary assets:
a clinical study of inter-firm and intra-firm mechanisms**

Abstract

This paper reports the results of a 4-year longitudinal clinical study conducted at executive office levels in Intel Corporation. It seeks to remedy the neglect of firm-level empirical analyses of capital budgeting, and in particular to address the mechanisms used to coordinate investment decisions and associated expectations. The aim is to provide an empirical illustration of recent work that has modelled formally the benefits available when capital spending decisions are structured as complementary investments at both intra-firm and inter-firm levels. Within Intel's capital budgeting process, we focus on a hitherto neglected mechanism termed a technology roadmap – a mechanism used to ensure that large-scale capital investments made by sub-units of the firm are coordinated with one another, and that they are aligned also with investments made by a wide range of complementor firms, such as OEM customers and developers of operating systems. We describe the technology roadmap mechanism, and examine how it integrates with DCF analyses to permit an individual capital spending proposal to be valued within the system of complementary investments of which it is a part. The contributions of the paper are threefold. First, our findings provide strong, firm-level evidence supporting the arguments of Milgrom and Roberts (1995a, b) and Trigeorgis (1995, 1996) that the system of assets, rather than the individual investment decision, may often be the critical unit of analysis and decision for managers. Second, we find that value-maximising investments in systems of complementary assets require coordination mechanisms that are largely overlooked in recent theoretical literature. Third, we identify issues for investigation in future large-sample surveys and clinical analyses of the capital budgeting process. In particular, we suggest investigating whether there are systematic differences between industries in the effectiveness with which interdependent investments are planned and coordinated across firm boundaries.

1. Introduction

The importance of coordinating individual capital investment decisions, so that the system of assets becomes the unit of analysis and valuation, is significant for firms and has been addressed in several recent studies. Focusing primarily on intra-firm coordination, Milgrom and Roberts (1995a, 1995b) have modelled the synergies between capital spending decisions as involving complementarity relations, such that additional investment in any one component of a system increases the returns to additional investment in the others. Under these conditions, value-maximizing results may be achieved only by coordinated change in all components, and not by altering one in isolation from shifts in the others. Trigeorgis (1996, p. 256) has commented that "skillful managers value the presence of interactions among various projects", and may be expected to incorporate synergies between parallel investments and between projects over time in their appraisal practices. These concerns with the coordination of complementary investments have been extended to the inter-firm level. Brennan and Trigeorgis (2000, p.4) have argued that the cash flows from an investment project are influenced not only by agents within the firm, but also by the actions of agents outside the firm, such as competitors and suppliers, and that these actions can in turn be influenced by – as well as influence – the actions of the agents within the firm. Others have argued in similar terms. Using a game-theoretic approach, Farrell and Saloner (1988) have examined the relative merits of committees, markets, and a hybrid comprised of both, as mechanisms for establishing inter-firm compatibility standards for component designs. Katz (1995) has compared joint ventures with market transactions and mergers, as means of assembling complementary inputs to research and development projects. And Farell and Katz (2000) have examined the incentives to innovate when two strictly complementary components are produced by different firms, one of which is a monopolist.

Despite such formal modeling of systems of complementary assets, there remains an empirical deficit in the study of the actual investment appraisal and coordination processes of firms (Jensen, 1993; Graham and Harvey, 2002). There is still surprisingly little systematic study of how capital investment decisions are actually made in practice (Jensen, 1993, p. 870). This is particularly so with respect to field-based or clinical research that looks intensively at the investment appraisal practices of a single firm using a wide variety of data. More specifically, little is known about the mechanisms (other than competitive markets) through which the coordination of investments and related expectations within and among firms is

achieved (Miller and O'Leary, 1997). Also, little is known about how the coordination mechanisms used by firms relate to financial evaluation techniques, such as NPV, payback and ROI, that form the core of traditional capital budgeting practices. While Graham and Harvey's (2001) recent survey of capital budgeting polls a large set of firms, poses a broad range of questions concerning whether and when particular valuation techniques are used, and provides unique information on the financing policies of firms, issues of investment coordination are not addressed specifically. For instance, their questionnaire does not ask whether managers consider the scope of an investment decision, what mechanisms enable them to define this scope and, if there are complementarities to be economized upon, what practices are used to coordinate investments within and among firms and to value the set of synergistic assets.

This paper seeks to remedy the neglect of firm-level empirical analyses of capital budgeting, and of the mechanisms used to coordinate investment decisions and associated expectations in particular, by reporting the results of a 4-year longitudinal clinical study conducted at executive office levels in Intel Corporation. More generally, the paper provides an in-depth analysis of an important issue for researchers in financial economics. It examines the mechanisms that may allow firms to derive the benefits indicated by recent models of complementary investments, and that might be difficult to identify and describe other than through clinical research (Jensen et. al., 1989; Tufano, 2001). We selected the particular firm and its industry because they are ones in which extensive systems of investment are coordinated on a frequently recurring basis. Within Intel's capital budgeting process, we focus on a hitherto neglected mechanism termed a technology roadmap. This is used to ensure that large-scale capital investments made by sub-units of the firm (in assets such as new processes, microprocessor products and manufacturing capacity) are coordinated with one another, and that they are aligned, also, with investments in enabling and related technologies on the part of a wide range of other firms, including those in Intel's supplier base, its OEM customers, and developers of operating systems, software and communication infrastructures. We describe the technology roadmap mechanism, and we examine how it integrates with DCF analyses to permit an individual capital spending proposal, in such as a new microprocessor product, to be valued within the system of complementary investments of which it is a part. We examine also the role of industry-level technology roadmaps produced by the SEMATECH consortium, and how these support firm-level coordination of investments and related expectations. By thus analysing the complexities of designing an

intra- and inter-firm coordination mechanism, we find support for particular lines of theoretical enquiry, and identify specific issues to be addressed in future clinical studies and in further refinements of large-sample surveys of the capital budgeting process.

The remainder of the paper is organized as follows. Section 2 describes our field research methods. Section 3 analyses the structure of the complementarity relations available to Intel. Section 4 examines the roles of technology roadmaps in coordinating investments at inter- and intra-firm levels. In section 5, we provide implications for future research and conclusions.

2. Methodology

Permission to undertake research within Intel Corporation was sought initially in negotiations with an executive vice-president of the firm. In light of the sensitive issues being addressed, approval was granted only subject to signing a formal, non-disclosure agreement. This allowed the researchers to gain access to private information, and to study the application of the firm's investment coordination and appraisal practices to a particular technology generation during the period May 1996 to June 2000. Release from the non-disclosure agreement was negotiated at the conclusion of the research, so that the firm's identity could be revealed. This process did not constrain the arguments and evidence presented in this manuscript, which do not depend on disclosure of confidential data relating to such as net present values for specific projects, or future product plans and strategies.

By negotiating access to the most senior managerial levels of Intel, and conducting a multi-year study, it was possible to identify sources of data and to examine materials relating to the firm's actual capital budgeting process that are inaccessible to survey-based and large sample studies (Wruck and Jensen, 1994; Graham and Harvey, 2001). Such a detailed and extensive piece of clinical research is unusual in the literature. However, any such study has the inherent limits of a small sample, with the inevitable constraint that its results may be sample specific. This may be overcome in subsequent research, in particular by utilizing the detailed empirical description provided for theory development and communication (Tufano, 2001).

Four research methods were used to compile a substantive data base. These were: interviews with key decision-makers; the manual collection and analysis of internal documents; first-hand observation of processes; and the collection and analysis of the public-record concerning the firm and the industry.

Given the concern to study the coordination of major capital investments, interviews were sought with many of the firm's most senior officers. Interviews were requested with 33 executives and managers, selected for their roles in making investment decisions and in developing and extending the firm's capital budgeting practices. All of those approached agreed to be interviewed. All interviews were conducted by the authors. Most of the interviews were held at Intel's corporate offices in Santa Clara, CA., and at its facilities in Chandler, AZ., Albuquerque, NM., and Hillsboro. OR.; the remainder were held at the firm's manufacturing facility in Leixlip, Ireland. Those interviewed included: the president and chief executive officer; the chief financial officer; vice-presidents for technology development, manufacturing, microprocessor product design, and marketing; the director of technology strategy; and managers and engineers in research and development facilities and high-volume factories. In addition, interviews were held with three technical analysts who focus exclusively on examining the semiconductor industry for the primary trade publications. They were asked to describe their understanding of Intel's coordination practices. All interviews were semi-structured and lasted a minimum of one hour. All but three of the interviews were tape-recorded.

The researchers gained access to and analyzed a range of documents confidential to Intel Corporation. These included the firm's capital investment manual, engineering and technical manuals, and the proceedings of intra-firm conferences that describe how investment appraisal and coordination practices were devised and how they have been modified and extended in use. Intel fabrication facilities in Ocotillo, AZ, Rio Rancho, NM, and Leixlip, Ireland, were visited, to gain a first-hand understanding of the firm's technology development and manufacturing processes.

Internal data sources were complemented by analyses of the public record concerning the firm and the industry. Press releases and press coverage were studied, as well as speeches by Intel executives, the proceedings of trade conferences, technical and trade journals, and the reports of technical and financial analysts.

3. The firm and its complementarity structure

Intel designs and manufactures microprocessors, the logic devices that enable computers to execute instructions.[1] Throughout the 1990s, its share of the world-wide market for PC microprocessors exceeded 80% of units shipped. During the same period, the firm's ratios of gross-profit and operating-profit to net revenues generally exceeded 50% and 30% respectively. The ratio of operating profit to total assets generally exceeded 20% (Figure 1), such that key analysts ranked Intel the world's most profitable microprocessor producer.[2] A key element in the firm's strategy has been to invest, at frequent intervals and in a coordinated manner, in improved fabrication processes, new products, and enhanced manufacturing practices.

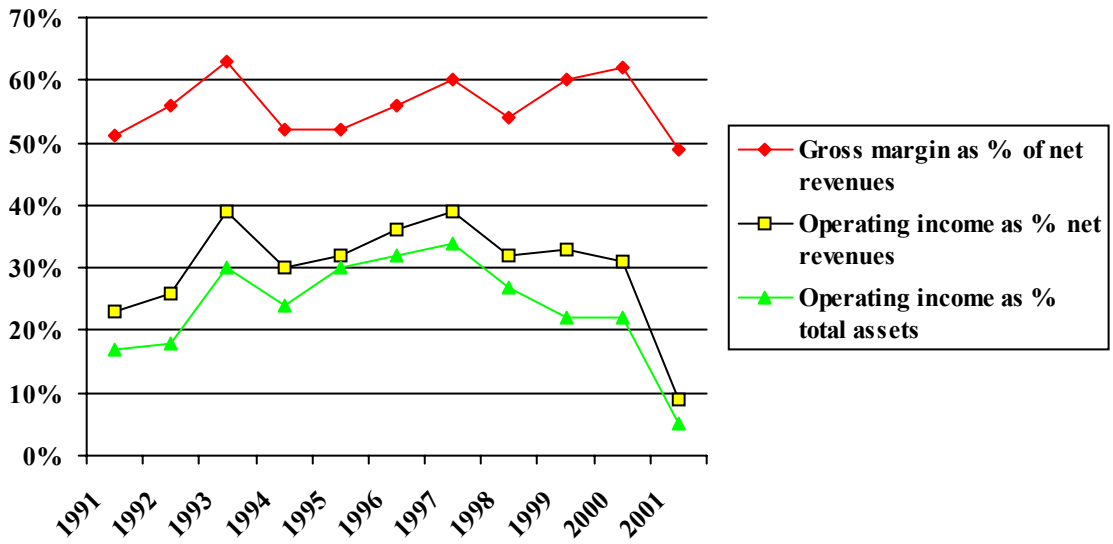

Fig. 1. Intel financial ratios, 1991 – 2001, derived from summarized financial statements of the Corporation reproduced in Appendix A

[1] The firm also manufactures hardware and software products for internet-based and local-area networking, as well as chip-sets, motherboards, flash-memories and other "building blocks" for computing and internet-based communication.

[2] M. Slater. Profits elude Intel's competitors. *Microprocessor Report* May 10, 1999. The recession that began during 2001 resulted in a decline in Intel's sales volume, and significant disimprovement in key operating ratios. It is unclear at this time whether, or how quickly, the firm will resume its historical levels of performance.

Since the mid 1980s, Intel has invested in an improved process for fabricating microprocessors, termed a process-generation, at intervals of ~3 years. In addition, and at comparable intervals, it has designed at least one new family of microprocessor products, and commenced manufacture in 3 – 6 geographically dispersed factories, each of them incorporating improvements in layout, operating policies, training and other procedures. This process of recurrent investment in both products and processes requires substantial levels of intra-firm and inter-firm coordination. Developers of Intel's proprietary process-generations collaborate closely with a range of suppliers such as Silicon Valley Group and Nikon, that are investing concurrently to design more advanced equipment-sets and materials. Without corresponding advances in lithographic equipment-sets by those firms occurring at defined moments, Intel would be unable to operationalize its successive generations of process technologies. The value of advances in microprocessor design would thus be substantially reduced. Also, Intel's microprocessor architects seek to coordinate their designs with those of customers and firms that are investing in complementary products. These include computing-devices by Dell, Compaq, Fujistu and others, operating systems by developers such as Microsoft and Linux, data-base management systems, and extensive sets of application-software programs devised by hundreds of firms around the world. Again, without these complementary investments being made by other firms, and their timing being carefully and accurately synchronized, the financial value to Intel of improvements in the speed of microprocessors arising from process and product advances would be substantially less.

Through the coordination of investments within the firm, and with both upstream and downstream firms, Intel's executives seek to economize on what Milgrom and Roberts (1995b, p. 200) have termed a "complementarity structure". In this section, we set out the components of this complementarity structure, as a prelude to examining in section 4 the mechanisms that are used to coordinate them. In the three sub-sections that follow, we examine the separate sets of relations comprising that structure. Firstly, we examine how they may arise when a new process generation is developed and operationalized concurrently with new microprocessor products. Secondly, we look at the benefits available when new microprocessor product designs align with complementary computing, operating system and software products. Thirdly, we consider how complements may be achieved when a new process-generation is accompanied by advances in the designs of Intel's high-volume factories. To illustrate the importance of successful

coordination, and how critical timing is, the fourth and final sub-section illustrates the costs to the firm of failing to align successfully the overall set of complementary assets.

3.1. Coordinated process-generation and microprocessor designs

The aim of investing in each new process-generation is to reduce the minimum linear feature-size of an electronic element, such as a transistor, so that more of them can be formed on a silicon wafer.[3] This increase in transistor density produces two main effects. First, it increases the yield of good microprocessor die per silicon wafer (die-yield). Second, it improves the speed at which a microprocessor can execute instructions (clock-speed).[4]

Intel's executives seek to establish and optimize complementarity relations by coordinating incremental investments in the manufacturing process that increase transistor density, and incremental investments in new products. The design of a new product generally consists of extensions to an architecture, so that the microprocessor can execute an enhanced set of functions at a faster clock-speed. A typical effect is to increase the number of electronic elements on the microprocessor die, thus increasing its area and reducing die-yield per wafer on a given fabrication process (see Appendix B). The returns to coordinated introduction of a new process generation and a new microprocessor are generally higher than to both changes made independently. The increased transistor density of the process at least partially offsets the larger die-size of the product[5], resulting in lower unit costs of manufacture. It also boosts the clock-speed increases that are achieved by improvements to the product architecture.

The coordination of investment in process-generation and microprocessor design forms the initial step in the production of complementarity relations. A second step is to seek to align the designs of the

[3] At present, electronic elements of 0.13 micron in length are being patterned on wafers and, historically, the length has been reducing by a factor of ~0.7 per process generation. A micron equals 1/1,000,000 of a meter.

[4] As feature-sizes are reduced, electrons take less time to complete an electronic circuit, thus enhancing the clock-speed of the microprocessor.

[5] The larger die-size of a new microprocessor could make its manufacture unprofitable, if it were planned and launched independently of a shift in fabrication process. For example, the die-size of Intel's 64-bit, Itanium, microprocessor was so "bloated" that manufacturing it on the 0.25-micron process during 1998 or 1999 would have produced significantly negative returns (L.Gwennap. Intel's two-track strategy re-routed. *Microprocessor* Report, August 4, 1997). Only by optimizing the product's design for the later, 0.18-micron process, and launching it when that process was in its mature phase, did the investment in Itanium indicate positive returns.

microprocessor products with those of complementary products such as hardware, operating systems and software, devised by other firms, so as to form markets for more advanced, end-user computing devices.

3.2. Coordinated microprocessor and complementary product designs

Intel's policy is to lead competitors in introducing new microprocesor products, and to coordinate the launch of each one with the introduction of more advanced computing devices, operating systems and application software designed by other firms:

> "My worry, right now, is to assure two things: one, that our strategies are aligned with our complementors, if you will; and second, is to accelerate their programs such that, when their product gets to the market, it is pretty much in-time with our product, not a year or two years later, so that can accelerate the volume deployment of a new generation of technology."[6]

Intel's share of the PC microprocessor market has been high since the beginning of the 1990s, generally exceeding 80% of world-wide unit shipments (Figure 3). Its revenue growth rate has come to depend upon forming and expanding markets for a succession of more advanced, end-user computing devices, each of them incorporating more powerful Intel microprocessors as core, logic devices:

> "[What] we started figuring out and I think implemented in the early 1990s was that, in fact, Intel had moved into a position, in terms of our market share and the particular products that we sold, that to try from a concentrated standpoint to take the other guy's market share from him wasn't going to sustain growth for us. .. So, we started moving into a mentality that went along the lines: if we can do things that stimulate the market growth, we will assume that we are going to take our fair share of that position."[7]

[6] Interview, Executive Board Member and President of Intel Capital, July 28, 1998.
[7] Interview, Manager, Technical Analyst Relations, August 24, 1998.

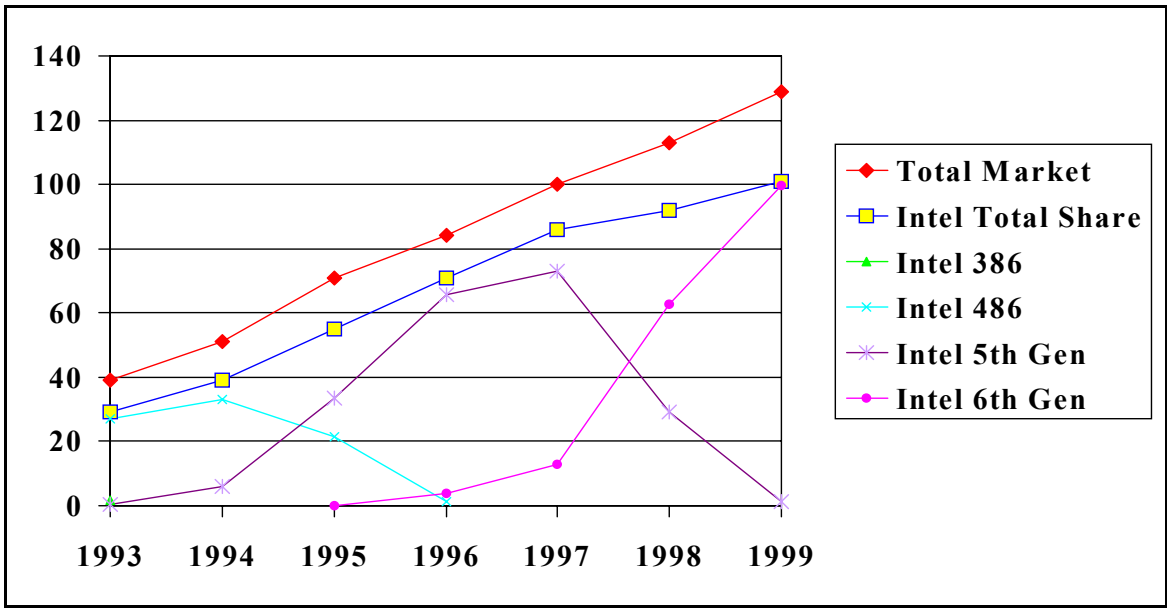

Fig. 3. Total global shipments of PC microprocessors, 1993 – 1999, and Intel's market share. Data are in millions of units.[8]

From this dominant postion within the microprocessor market, Intel aims to produce complementarities that are available through coordinating investments at the inter-firm level. The timing of the launch of a new microprocessor is critical, since Intel usually introduces a new microprocessor at a relatively high price, which is then reduced significantly during the product's short life cycle. The aim is to secure product acceptance on the part of the most demanding users initially, while the product is still manufactured in low volumes in the development factory, and then to stimulate demand growth by lowering prices as additional factories are brought on-stream. Life cycle revenue is thus significantly higher for Intel when its product investments are coordinated successfully and precisely with those of related firms, such that a new microprocessor, enhanced operating systems, improved Internet infrastructures, and novel software applications are all available from the outset of a given generation.

[8] Gwennap and Thomsen, Intel Microprocessor Forecast (Sebastopol, CA: Micro Design Resources Inc., 1998).

3.3. Coordinated process-generation and factory designs

The third element in the complementarity structure involves the coordination of investment in each process-generation with investment to enhance Intel's high-volume manufacturing capabilities.

While successive process generations offer increases in die-yield and clock-speed, each one also involves working to finer tolerances, across a greater number of manufacturing steps, using several equipment types and materials that are new to the firm and to the industry. Performance levels achieved in the development factory become more difficult to sustain as successive process-generations are transferred to high-volume manufacturing facilities, whose personnel have to learn the parameters of increasingly complex systems. Lower performance levels during the learning period could require investment in excess capacity to achieve a given level of output, thus diminishing the benefits Intel gains from stimulating high-priced, early-period demand for new microprocessors.[9]

The firm seeks complementarities by coordinating the introduction of each process generation, offering enhanced die-yields and clock-speeds, with advances in factory design aimed at reducing the time to learn new system parameters. Since the early 1990s, and to combat the so-called "Intel-u"[10], the firm has sought closer integration of its development site and high-volume factories, using "virtual factory" control practices. The intent has been to engineer each generation of high-volume factories so that it more closely copies and reflects the exact layouts, equipment sets, operating procedures and intervention policies established in the development site. The trajectory of improved performance in the development site is thus continued within each of the high-volume factories, as though the network as a whole comprised a single manufacturing entity.

3.4. Costs of a coordination failure

There are costs of coordinating investments in process, product and factory designs with one another internally, and with those of suppliers, complementors and customers externally. They include the expense

[9] Interview, Director of Technology Strategy, December 11, 1996.

of the organization structures and systems by which various groups align their design decisions. Also, there are costs of rendering product development resources fungible, so that, for instance, groups of architects may be re-assigned to develop a particular microprocessor more quickly to synchronize with the earlier availability of a process-generation. Historically, Intel executives have found such expense to be substantially lower than the benefits:

> "We will take a new process [generation] as soon as we can get one, and we will put as many products on the new process as we can, and incur any [incremental] cost necessary. The return is so great on moving to the next generation that it's the [process] technology that's the "preventer". It's the ability of actually being able to engineer the next process that's the "limiter".[11]

Table 1 estimates the manufacturing costs of one hypothetical coordination failure, in which the 0.25-micron process generation becomes available one quarter later than the Pentium II microprocessor product. It is assumed that volume of sales for the quarter remains unchanged, but in the absence of the newer fabrication technology the Pentium II would continue to be manufactured on the earlier, 0.35-micron process generation. As a consequence, the product's die-size is larger and the yield of good die is lower. Each wafer produces only 58 good die, compared to 120 if the newer fabrication process were available. Although fabrication costs are higher for the 0.25-micron process, the net effect of the delay is excess manufacturing cost of $480 million, almost 6% of Intel's operating income for the relevant year, 1998.[12] Even relatively short lags between the arrival of a fabrication process and a product may thus result in significant diminution in Intel's operating income.

---

[10] The phrase is part of Intel folklore. It refers to the early history of process transfers, when product yield would decline significantly each time a process generation was transferred from development to high-volume factories, and would remain depressed for several months, resulting in a u-shaped yield curve.

[11] Interview, Chief Financial Officer, Intel Corporation, August 26, 1998.

[12] If the process-generation became available one–quarter *earlier* than the product, Intel would have encountered the opportunity costs of marketing microprocessors of lower clock-speed that required greater areas of silicon die to manufacture. Data to estimate such costs are not available from the firm or from analyst reports.

Table 1. Estimated manufacturing cost of a failure to coordinate process-generation and product designs[13]

Condition:	Process lags product by 3 months	Synchronized designs
Process-Generation:	0.35 micron	0.25 micron
Product:	Pentium II	Pentium II
Die-size & yield data:		
Microprocessor die-size (mm^2):	$203mm^2$	$131mm^2$
Yield of good die per silicon wafer:	58	120
Estimated manufacturing costs per good die:		
Fabrication: Package: Packaging and testing: Module parts and assembly:	\$49 \$16 \$15 \$14	\$28 \$16 \$12 \$14
Total manufacturing cost per good die:	\$94	\$70
Manufacturing cost of coordination failure		
Unit cost difference [\$94 - \$70]	\$24	
Volume (Q1, 1998 est. unit shipments of Pentium II):	20 million	
Estimated total cost of coordination failure:	\$480 million	
Excess cost as % 1998 operating income (\$8,379,000,000):	5.7%	

In the following section, we analyze how Intel seeks to avoid such costs, and to realize the benefits available from the complementarity structure, through practices of intra-firm and inter-firm investment coordination.

4. Technology roadmaps

Consistent with the large-scale firms surveyed by Graham and Harvey (2001), Intel's capital budgeting process requires discounted cash flow analyses. Net present values are calculated for proposed new microprocessors within the product development groups, for instance.[14] Net present cost analyses are used extensively, as when factory planners are choosing between capacity installation alternatives, such as

[13] Intel Corp., Microprocessor Reference Guide (2000) and press releases; L. Gwennap and M. Thomsen, Intel Microprocessor Forecast (Sebastopol, CA: Micro Design Resources Inc., 1998).
[14] Interview, Vice President, Microprocessor Products Group, July 25, 1996.

whether to re-fit an existing facility for a new process generation or build from a green-field site, or whether to expand production in one country rather than another.[15]

In light of the extensive set of complementarities available to the firm, however, the capital budgeting process restricts the right of sub-units to evaluate investments "independently at each of several margins", in Milgrom and Roberts' (1990, p. 513) phrase. To be approved, an investment proposal must not only promise a positive return. It must be shown also to align with a technology roadmap.[16]

A technology roadmap sets out the shared expectations of the various groups that invest to design components, as to when these will be available, and how they will inter-operate technically and economically, to achieve system-wide innovation. Typically, a technology roadmap will address each of several future coordination points, defined by a year or quarter-year. The groups involved in preparing it may include sub-units of a firm, as well as suppliers, complementors and OEM customers. A roadmap is an inherently tentative and revisable agreement, one of whose key roles is to enable design groups to assess the system-level implications of advances, delays or difficulties in bringing investments in new component designs to fruition.[17] Equally, the expectations reflected in a technology roadmap may require fundamental revision if there are indications of insufficient demand for the end-user products to which the system of component innovations is expected to give rise. A roadmap thus provides a mechanism for the dynamic coordination of expectations where there is recurrent, intra-firm and inter-firm investment.

Through linking an investment explicitly with a technology roadmap, the proponent is required to demonstrate that it synchronizes and fits with other, related investments, both within and beyond the firm, in ways that promise to maximize returns. Ensuring that individual investment decisions are congruent with the relevant roadmap is afforded the highest priority by Intel's executive officers. Such is the importance of the complementarity structure for the overall profitability of the firm that the CEO addresses the coordination of decisions directly:

---

[15] Interview, Chief Financial Officer, Intel Corporation, August 26, 1998. Net present cost analyses establish discounted cost differentials, taking revenue to be the same across alternatives.

[16] Intel Corporate Finance, Capital Project Authorization (1998) (internal document); Interview, Corporate Capital Controller, July 23, 1996.

[17] However, the costs of revision to individual sub-units and firms may increase as a particular coordination node approaches, because each will have invested in the expectation of system-wide success.

> "We obviously do ROIs on products and things of that sort, but the core decisions the company makes, the core decisions are basically technology roadmap decisions, driven off-of a fundamental belief that Moore's Law will continue to be valid, that the purchasing community will continue to buy new [computing devices], and that our main charter is to stay ahead of our competition in that space."[18]

In the sub-sections that follow, we analyse and illustrate how a technology roadmap is prepared and the roles it plays in investment coordination. We follow the chronology of roadmap preparation, beginning with the alignment of investment decisions between Intel and firms in its supplier base. Our argument is that the active management of investment programs (Brennan and Trigeorgis, 2000) may depend upon such a mechanism, so that decisions to make, defer or otherwise alter a particular capital commitment may be made in the light of related shifts in plans and expectations elsewhere, throughout a system of investments being made at intra- and inter-firm levels.

4.1 Coordination with suppliers' innovations

Intel depends upon innovations by suppliers of equipment-sets and materials to operationalize each of its new process-generations, and thus begin its cycles of complementary investment in process, product and factory designs. The firm regards such innovations on the part of suppliers as benefiting the industry as a whole, and cooperates with other semiconductor manufacturers to specify collective design needs and time-lines:

> "[It's] much more economical for our industry to work as a whole to create some base technology, and the real intellectual property, the real value-added, comes not from creating a stand-alone piece of lithographic equipment, or a stand-alone piece of ion implanter [equipment]; it comes from the integration of those into a total process. So our intellectual property, or trade secrets in this case, comes from the integration. We're very able to work with our competitors in creating the stand-alone pieces."[19]

Coordination of investments by semiconductor firms and their supplier base is facilitated by a technology roadmap that is prepared under the auspices of the SEMATECH consortium. Table 2 shows

[18] Interview, President and CEO, Intel Corporation, December 17, 1998. By "ROIs", the CEO means summary financial statistics, including net present value and net present cost, as mandated by Intel's Capital Project Authorization manual. "Moore's Law" is named for Intel co-founder and chairman-emeritus Gordon Moore, who noted in 1975, and on the basis of empirical observations extending across 15 years, that the semiconductor industry seemed capable of doubling the number of electronic elements on a memory device every 18 months, at no increase in cost per device. See Moore (1975).

top-level statistics from such a roadmap that was published in 1994. It was prepared by delegates from each of the 13 firms comprising the consortium, including Intel, which accounted collectively for over 80% of US output of semiconductor devices. They collaborated with the trade association representing supplier firms through joint working groups and conferences, and liaised also with relevant US federal and university laboratories. The resultant roadmap indicated the design requirements for equipment-sets and materials at each of five future coordination points. Its preparation may be divided for analytic purposes into three steps.

The first step was to specify rates and directions of change in individual design variables to achieve coordinated results at each point. The aim was to stimulate and inform suppliers' investments in research and development and in the commercialisation of new equipment sets and materials, by indicating when the US semiconductor industry as a whole would require equipment sets and materials of particular tolerances and capabilities, and in sufficient quantities for high-volume manufacture. The changes in design variables were specified by extrapolation from historical performance levels, specifically, by assuming that the innovative conditions under which Moore's Law had been achieved in the past could be made to persist:

> "[It's] a self-fulfilling prophecy. Moore's Law is not a law of physics. On the other hand, it's a pretty strong economic law, because once the industry deviates from Moore's Law, then the rate of investment is going to change, and the whole structure will change. And at the point we deviate from it, it means the industry is maturing. [Moore's Law] just turned out to be possible early-on, and then it got built-into the economics of things."[20]

Reduction in minimum electronic feature-size at a constant rate of 0.7 per coordination point, due to investments in innovation by lithography suppliers, was expected to continue through the life-time of the roadmap, and to combine with anticipated rates of increase in wafer diameter achieved by suppliers that develop silicon (Table 2). Coordinated availability of these and other components was to permit semiconductor firms to continue to operationalize new process generations, that would increase the number of bits on a memory product by a factor of four[21], and the number of transistors on a microprocessor die by a multiple of ~2.3.

---

[19] Interview, President and CEO, Intel Corporation, December 17, 1998.
[20] Interview, Manager of Lithography Process Equipment Development, November 3, 1997.
[21] This is the rate of increase in electronic elements on a memory device that Moore's Law calls for, viz., a multiple of 4 per 3 years, or 2 per 18-month period (Moore, 1975). The industry established a different

Table 2. Required rates and directions of change in individual design variables to achieve coordinated & system-wide innovation. As specified in National Technology Roadmap for Semiconductors (1994)[22]

	Current	Future				
Technology node	(N_0) 1995	(N_1) 1998	(N_2) 2001	(N_3) 2004	(N_4) 2007	(N_5) 2010
Suppliers' innovations in equipment-sets and materials[23]						
Lithography						
Minimum feature size (microns)	0.35	0.25	0.18	0.13	0.10	0.07
Scaling factor per generation		*~0.7*	*~0.7*	*~0.7*	*~0.7*	*~0.7*
Silicon wafers						
Wafer diameter (mm)	200	200	300	300	400	400
Increase per two generations (mm)			*100*		*100*	
Advances in semiconductor product designs:						
Memories						
Bits per die (millions)	64	256	1000	4000	16000	64000
Multiple per generation		*4*	*~4*	*4*	*4*	*4*
Cost/bit (thousands of a cent)	0.017	0.007	0.003	0.001	0.0005	0.0002
Scaling/ reduction factor		*~0.45*	*0.5*	*~0.5*	*0.5*	*~0.5*
Microprocessors						
Transistors per die (millions)	12	28	64	150	350	800
Multiple		*~2.3*	*~2.3*	*~2.3*	*~2.3*	*~2.3*
Cost/transistor (thousands of a cent)	1	0.5	0.2	0.1	0.05	0.02
Scaling/ reduction factor		*0.5*	*~0.5*	*0.5*	*0.5*	*~0.5*

The second step was to seek an industry wide consensus on which research and development approaches were most likely to provide the enhanced materials and equipment-sets required at each coordination point. The aim here was to guide the sets of investment alternatives to be considered and analysed by suppliers. SEMATECH working groups compared needs with the anticipated state of R&D for individual types of equipment-sets and materials, at each point. In the case of the more immediate ones, such as 0.25- and 0.18-micron, the aim was to determine whether incremental improvements to existing technologies could be commercialised in time. In the case of later coordination points, it was to identify

---

constant for increases in microprocessor functionality, viz., a rise in the number of transistors per die by a multiple of ~2.3 every three years.

[22] Adapted from Semiconductor Industry Association, National Technology Roadmap for Semiconductors (San Jose, CA: SIA, 1994, p. B-2).

[23] For brevity of exposition, only two types of components whose designs are coordinated are included here; the full version of the roadmap includes many others, such as deposition and implantation equipment, mask technologies, etc.

the limits of extant technologies and, by providing an intensive, industry-wide assessment of the state of research and development, to focus the attention and the investments of often smaller suppliers and laboratories on the most promising alternatives. Thus, for instance, competition was encouraged between suppliers to develop and commercialise extreme ultraviolet, e-beam projection and proximity x-ray technologies as replacements for deep ultraviolet lithography, for patterning transistors on silicon at feature-sizes of 0.1-micron and below.

The third and final step was for individual firms to decide whether to align their investment programs with the roadmap, whether to seek modifications to it, or whether to withdraw from the consortium. To date, Intel executives have found extensive net benefits from participating in SEMATECH, not least the ability to share component design costs with other semiconductor firms, and to increase suppliers' incentives to invest by specifying when industry-wide demand should commence for components of particular capability. However, the CFO does require periodic appraisals of whether the firm would benefit from negotiating modifications to the roadmap.[24] During 1994, for instance, Intel executives concluded that 2-year innovation cycles were more likely to be optimal for the firm than the historical, 3-year duration. The decision was based on a discounted cash flow analysis of whether more frequent increments in transistor-density and microprocessor clock-speed, available from 2-year cycles, would outweigh such costs as faster process-generation and product obsolescence. In extensive negotiations with consortium members and the supply industry, a temporary shift to 2-year cycles was agreed with respect to the 0.25-, 0.18- and 0.13-micron nodes, with a reversion to 3-year cycles thereafter (Table 2).[25]

The SEMATECH technology roadmap thus provides a mechanism for coordinating expectations and investments among a set of firms and its supplier base in a key sector of the modern economy, where there is recurrent and system-wide innovation. In addressing design requirements comprehensively for all core types of components, it reflects the dependence of investment returns to any one specialized firm on close coordination with the design plans of others:

> "In semiconductor manufacture, progress tends to occur in discrete generations where all the technology elements need to be in place before a transition can be made to the next generation."[26]

---

[24] Interview, Chief Financial Officer, Intel Corporation, August 26, 1998.
[25] A revised version of the SEMATECH roadmap incorporating the changes was published during 1997.
[26] Semiconductor Industry Association, National Technology Roadmap for Semiconductors (San Jose, CA: SIA, 1994, p. 27).

Partial coordination of a system of investments may not come close to producing optimal returns in this industry, an observation consistent with the implication that Milgrom and Roberts (1995b) derive from their models of complementarity relations. By establishing where design lags are most likely to occur at each of several future nodes, and then identifying and monitoring promising lines of research, the technology roadmap provides individual firms that are making investments with a reasoned and expert assessment of the availability of complementary components. And by affording opportunity to lobby for changes in the roadmap, the SEMATECH process acknowledges the inherently high levels of uncertainty affecting all parties, and the need to focus attention and resources on any unexpected technical and financial difficulties affecting particular firms or sectors.

In this last respect, the SEMATECH process helps to focus the venture capital spending of a large and profitable firm, such as Intel, on "repairing" the roadmap so as to protect the complementarities available from investments within its core areas of business.[27] Early in 1995, for instance, executives from Intel, Texas Instruments and Motorola concluded that a particular way of advancing deep ultraviolet lithography, one reflected in the step-and-scan technology of Silicon Valley Group, would meet their collective needs for the 0.25-, 0.18- and possibly 0.13-micron nodes. To ensure the commercialisation of the component and its availability in sufficient volumes, the three firms took a joint equity stake in Silicon Valley Group, a relatively small and low-margin supplier.[28] During 1997, Intel, AMD and Motorola formed a private industry consortim to advance extreme ultraviolet (EUV) lithography, as an alternative for patterning feature-sizes on silicon at 0.1-micron and below. The consortium invested $250m in EUV research programs at three US Department of Defence laboratories.

4.2 Intra-firm coordination

In light of the shared expectations formed with suppliers, Intel managers continue the roadmap preparation procedure inside the firm. They plan several future process generations to coincide with the availability of more advanced equipment sets and materials. Three primary pieces of data are recorded in the

[27] Interview, President of Intel Capital, July 28, 1998.

intra-firm roadmap with respect to each generation: when it is expected to be available for test production and high volume manufacture; the key technical changes it is to introduce, particularly with respect to additional transistor density; and the expected capital investment to install a unit of capacity utilising the new process. [29] The data is communicated to Intel's factory design group and microprocessor architects, so that they may extend the intra-firm roadmap to show the combined financial effects of aligning the introduction of each process generation with that of more advanced manufacturing practices and new products.

In 1994, for instance, the intra-firm roadmap showed the planned availability during 1997 of a process generation to pattern 0.25-micron transistors on silicon wafers (Figure 3). To partially offset the rise in investment per unit of capacity, associated with the more advanced process, factory designers sought to coordinate its introduction with that of improved manufacturing layouts and operating policies in the high volume factories:

> "I am designing policies hand-in-hand with the people who are currently developing [a process-generation]. So it is meant to be a continuum. .. [We] design a continuum of policies, so that we have a set of policies that's intended to maximize information turns in a technology development factory, and in early high-volume factory to maximize output, late high-volume to minimize cost, ramping to maximize the ramp velocity. We need - in a factory, at a given snapshot in time - a WIP policy, an equipment maintenance policy, a cross training policy, etc., etc., that fit together." [30]

Of particular concern was to increase ramp-velocity, by altering factory layouts and equipment installation, staffing, and operating policies. Ramp-velocity is a measure of how quickly a new process generation can be "copied" from its development site to high volume factories without impairing a given level of die-yield. The faster this is achieved, the lower the total investment needed to meet a given volume of demand, and the greater the financial benefits of a new process generation.

---

[28] Interview, Manager of Technical Analyst Relations, August 24, 1998; D. McGrath, "SVG gets $129m lith order", Electronic News, July 26, 1997.

[29] A unit of capacity is measured as a given number of wafers introduced into production in a week (e.g. 5,000 wafer-starts-per-week). Capital investment data is only communicated selectively within the firm, to senior managers who require it as input to their investment proposals.

[30] Interview, Principal Scientist, Manufacturing Systems, August 22, 1997.

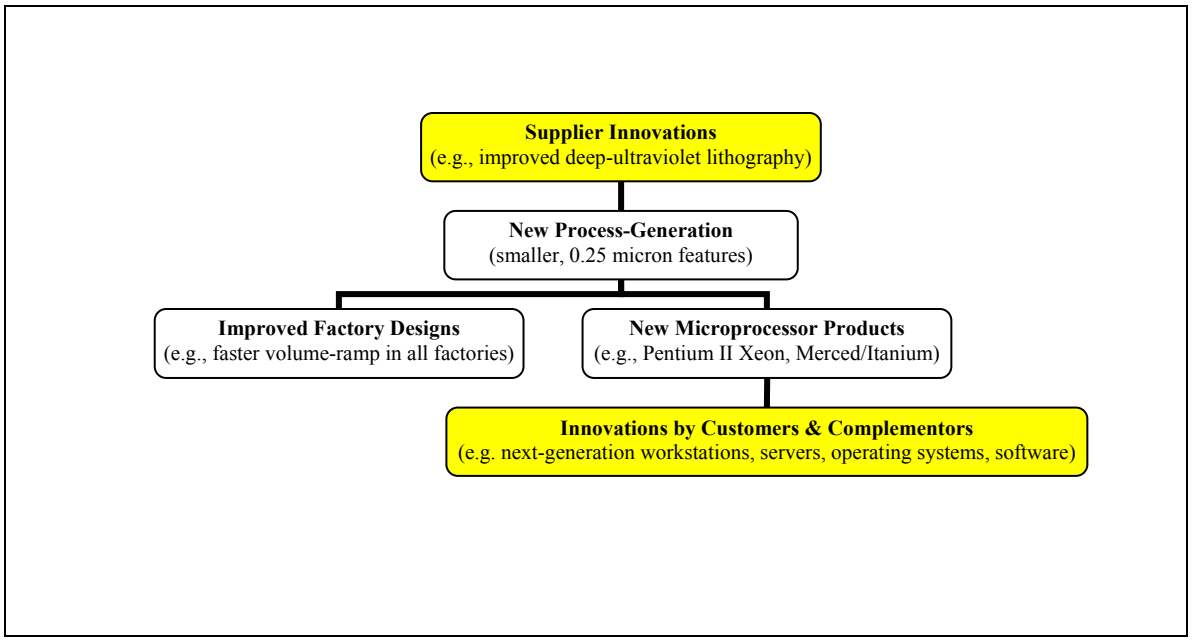

Fig.3. Components of the 0.25 micron technology generation whose design Intel sought to coordinate at intra- and inter-firm levels. Components developed by other firms are indicated by shaded boxes.

Microprocessor architects extended the intra-firm roadmap still further, by planning the investment schedules and time-lines of several new product families to coincide with the availability of the new process. By examining this alignment, we demonstrate the roles of a technology roadmap in permitting capital spending on new products to be appraised within the system of complementary assets of which they are to form a part.

Capital spending on a new microprocessor is typically proposed in stages, during a period of 4 or more years. Early investment is aimed at deriving a general model of the enhanced capabilities the new product might deliver for particular market segments, without commitment to a precise time-frame for execution or to manufacture on a given process-generation. But as architects move from that model to instantiating the new product as a set of circuits, layouts and masks necessary for manufacture, returns to additional investment come to depend significantly on coordinating product design closely with that of a particular process-generation:

> "If I set my [product] design target on the n^{th} technology [i.e., process-generation], and begin my implementation, begin my layout, I may spend on the order of a hundred-engineer-years of creating a physical layout only to find that I have to re-do it for the next generation [process] technology. So, the order of magnitude of investment, once I get to the point where I just want

> to instantiate transistors, is of the order of a hundred-[engineer]-years. And it might be several times that."[31]

A technology roadmap provides a mechanism for appraising whether such irreversible investment is justifiable in light of the investment time-lines and expected capabilities of complementary components.

During the early 1990s, for instance, Intel executives decided that, in addition to designing further products within its 32-bit architecture, the firm would also develop a line of new, 64-bit microprocessors aimed at higher-end workstation and server markets. A processor code-named Merced, devised jointly by Intel and H-P, was planned as the first instantiation of the new architecture. By consulting the technology roadmap, product architects sought to align their investment in the new product with the availability of a suitable, new process-generation:

> "[The technology development] organisation is very good at putting out a roadmap internally as to when they expect a certain process-generation to arrive. It is based on history of how often we have been able to increment the process-generations, and based on a forecast by some people in [the] organisation that are continually looking at where they expect, for example, lithography to evolve [by] a certain point of time. So, the [product] design group and myself, or general manager at the time, would have access to this technology roadmap .. that says, basically, as a function of time, this is the beginning point of the ramp of the .35-micron generation, for example, this is the entry point of the .25-micron generation, this is the entry point of the next generation that will follow that. .. The decision [on coordinating] a high-end product like this Merced [with a particular process-generation]… is actually very easy, in the sense that your product is oriented for performance. There is only one promise that you have [for customers] on this product, and that is that you'll offer the highest performance capability at the time for these high-end systems. So, you want to implement that on the most advanced [process] technology that would be available for manufacturing at the time the product would come out."[32]

The initial determination of the product architects was that the Merced should be introduced during the life-cycle of the 0.25-micron process-generation during 1998 or early 1999 (Figure 3). They believed that the product time-line could be made to align with that of the process, that the size of the product would permit an acceptable die-yield per wafer using transistors of 0.25-micron in length, and, generally, that an acceptable NPV would result from such a coordination.

The decision to launch a powerful and large die-sized product such as the Merced on the 0.25-micron process was based on a key assumption that the product would quickly be shifted to the newer, 0.18-micron process-generation. Not only was that generation expected to offer a further increment of transistor density, it was also anticipated that it would operate on larger, 300mm silicon wafers, which were in the

[31] Interview, Vice President, Microprocessor Products Group, July 25, 1996.

course of being developed by suppliers. As a consequence, the relatively large die-size of a product such as the powerful Merced product would quickly be offset by process-generation advances, such that an acceptable, long-run yield of good die per wafer could be achieved. However, unexpected revisions to the process roadmap in October 1997 led to a fundamental revision of such expectations:

> "The 0.25-micron process is in production here at Santa Clara, and it's coming up at other sites. The 0.18-micron process is in development. And, until recently, the plan was that 0.18-micron would become a 300mm generation [i.e., would operate on larger, 300mm wafers]. And we have just decided that it won't; it will stay 200mm. 300mm will shift out a generation. The reason is we were too fast. We have outrun the availability of equipment. .. That decision is two weeks old now, which is why you may not have heard of it."[33]

The expectation that suppliers would invest to devise the larger wafers, and supply them in high-volume in time for the 0.18-micron process generation, had proven to be incorrect. In addition, as the Merced's designers sought to perfect a novel, 64-bit architecture jointly with Hewlett Packard, they found during 1997 that the die-size of the product had become 'bloated' compared with initial expectations.[34] Such an unexpected revision could transform the economics of a product such as Merced. A key role of the technology roadmap mechanism is to convey such shifts in expectations, which may arise inside or outside the firm, to product developers to inform their capital investment decisions. Influenced by the delay in arrival of the larger wafer size, and also by difficulties in perfecting the Merced's instruction set, Intel's executive officers decided during 1997 to defer its launch, and the product's development time-line was reset so as to coincide with a later process-generation.

However, the time line and technical attributes of the 0.25-micron process was found to be fully aligned with those for a second family of new microprocessors, the Pentium II:

> "Pentium II was clearly the flagship product of our 0.25-micron technology. I want to make sure that the 0.25-micron technology is well suited for this product."[35]

This involved close collaboration between process engineers and product architects so that, as the Pentium II instruction set was refined and as its circuits and layouts were completed during 1996 and 1997, the

---

[32] Interview, Vice President, Microprocessor Products Group, July 25, 1996.
[33] Interview, Manager of Lithography Process Equipment Development, November 3, 1997.
[34] Interview, Chief Financial Officer, Intel Corporation, August 26, 1998; L. Gwennap. Intel's two-track strategy re-routed. *Microprocessor Report*, August 4, 1997. To correct for such unanticipated delays in completing any one microprocessor, Intel's policy is to design several new products in parallel design groups. Development of an alternative product may thus be accelerated through transfers of architectural skills and other resources, to protect the firm's competitive position in given market segments.

emerging 0.25-micron process-generation was adjusted to support features critical to its performance. Intel personnel thus sought to maximize the clock-speed of the new product while keeping its die-size sufficiently small for economic manufacture. The Pentium II contained 7.5 million transistors, 36% more than its direct predecessor, the Pentium Pro. But coordination of decisions on the part of product architects and process engineers resulted in a die-size for the new product that was actually 33% smaller than that of the Pentium Pro (Table 3). Also, whereas architectural improvements alone would have boosted the clock-speed of the Pentium II by ~50%, closely aligning its developmenta and that of the 0.25-micron process resulted in a speed increase of 125%. Complementarities are thus sought through coordinated product and process designs that combine improvements in clock-speed, which increase the marketability of product, with combined reductions in its die-size that reduce fabrication cost.

Table 3: Relative performance indicators for the Pentium II microprocessor[36]

Process-Generation			
Minimum feature-size (microns):	0.35	0.35	0.25
Products			
Brand name:	**Pentium Pro**	**Pentium II**	
Version:	Redesign	Original	Redesign
Date of 1st shipment:	2Q'96	2Q'97	4Q'97
Performance Indicators			
Die size			
Transistors per microprocessor (millions): Increase on Pentium Pro product (%):	5.5m	7.5m ~36%	7.5m
Microprocessor die-size (mm^2):	196mm^2	203mm^2	131mm^2
Die size *increase* due to architecture enhancement (%):		~4%	
Die size *reduction* due to process generation shift (%):			~35%
Die size *reduction* on joint product & process changes (%):			~33%
Clock-speed			
Maximum product clock-speed:	200MHz	300MHz	450MHz
Speed increment due to product architecture improvement (%):		50%	
Speed increment due to process generation shift (%):			50%
Speed increment on joint product & process changes (%):			125%

However, realizing the incipient benefits of new process and microprocessor generations depends on whether other firms devise more advanced, end-user computing devices, and markets for them, so as to

[35] Interview, General Manager, California Technology and Manufacturing, December 17, 1998.

accelerate the high-volume deployment of Intel's products. To that end, the firm's executives seek to ensure that their technology roadmap is aligned with those of OEM customers and complementors.

4.3 Coordination with customers' and complementors' designs

Since the early 1990s, Intel has taken a direct interest in the formation of end-markets for the varied types of products that incorporate its microprocessors. For instance, in the case of a particular version of the Pentium II, the Xeon processor, Intel coordinated its development with that of other firms' workstation and server computers, operating systems, data-base management systems, and an extensive range of applications software, in such areas as electronic commerce, supply chain management, and mechanical design automation. The aim was to ensure that these firms would invest to "integrate, tune, and optimise [their] solutions around this .. new microprocessor",[37] thus expanding Intel's market shares in the enterprise computing segment.

In seeking to align its plans with those of downstream firms, Intel shares elements of its technology roadmap with them, on a reciprocal basis and under non-disclosure agreements, for a period of up to two years prior to the planned product launch dates:

> "So, about the time that we are freezing on the product that we want to design, and looking forward to two years of design for its introduction, we have to take that to the software community and say "fine, here are the 70 new instructions that this processor has which will make [for example] your multi-media applications better", under non-disclosure agreement. "Here they are, start designing the product". So, [we take that data to] the software community, and the hardware community, and you also get the [technical analyst] people who make a living out of following our industry .. telling them "this is the direction that Intel's going in"". [38]

The sharing of roadmap data with technical analysts, thus going beyond the firms that are directly involved in product development, is integral to the coordination of investments at the inter-firm level. Bringing about complementary investments at the inter-firm level may depend on whether the parties have means of attesting the reliability of each others' claims and promises. In particular, smaller software vendors may be unwilling to invest if they lack confidence in the claims that Intel makes for its future microprocessor

---

[36] Intel Corp., Microprocessor Reference Guide (2000) and press releases; L. Gwennap and M. Thomsen, Intel Microprocessor Forecast (Sebastopol, CA: Micro Design Resources Inc., 1998).
[37] Intel Corporation press release, "Intel Pentium II Xeon processor launch", June 29, 1998.
[38] Interview, Chief Executive Officer, Intel Corporation, December 17, 1998.

generations. As one means of addressing such issues, Intel sometimes provides support in the form of technical assistance and venture capital to such firms.

But since about 1993, and also to assuage such concerns on the part of downstream firms, Intel has availed of the services of a small number of independent technical analyst firms, among them, Micro Design Resources:

> "The view we have come to have is that, in some way, we are the community organizer. We have brought together this community of people which cares about microprocessors. And we collect information from various of them, and then disseminate it out to all of them, and then we get them all together for meetings where they can all talk to each other. And so, I think you could say the biggest picture thing we do is we provide a focal point and opportunities for interaction of that community."[39]

Intel informs Micro Design Resources of key technical changes that it plans to incorporate in each of several future products, indicating also the particular market segment to which each one is being addressed, and its expected price point. The analyst firm's income stream depends significantly on the perceived objectivity and accuracy of its appraisals of such microprocessors on the part of customers who buy its newsletters, which include firms throughout the semiconductor, hardware and software industries, as well as stock analysts. Equally, Intel's willingness to continue sharing data with the analyst firm depends on the latter's adherence to product appraisals that, while they may on occasion be critical, nevertheless adhere to non-disclosure agreements with respect to proprietary data.

A technology roadmap thus provides a mechanism for the coordination of investment decisions throughout a design network, extending from suppliers to various sub-units within a firm and to its OEM customers and complementors. In the section that follows, we examine the implications of the mechanism that we have described for future large-sample studies and clinical analyses of the capital budgeting process.

5. Implications and conclusions

This paper has reported the results of a clinical study of how a major firm in the microprocessor industry coordinates and appraises investments in systems of complementary assets. It has described the overall complementarity structure within which Intel operates, both intra-firm and inter-firm, and demonstrated the costs of failing to coordinate successfully the sets of complementary assets. The role of

---

[39] Interview, President, Micro Design Resources, July 7, 1998.

technology roadmaps in coordinating both investments and expectations has been documented for the sub-units of Intel, and for the relations among Intel and its suppliers, complementors and OEM customers. The links between roadmaps as coordination mechanisms, and traditional capital budgeting practices, have also been analysed. We argue that the paper makes the following three contributions.

First, our findings provide strong, firm-level evidence supporting the arguments of Trigeorgis (1995, 1996) and of Milgrom and Roberts (1995a,b) that the system of asssets, rather than the individual investment decision, may often be the critical unit of analysis and decision for managers. In the case of Intel, analyzing "synergies among parallel projects undertaken simultaneously" (Trigeorgis, 1996, p. 257) is the aspect of investment appraisal that is always considered at the highest levels in the firm because, as we have demonstrated, the costs of failing to coordinate such complementary investments may be very high. By contrast, "interdependencies among projects over time" (ibid., p. 257), the traditional focus of real options studies, receive less formal attention at top-management levels. Our findings thus provide support for the extension of theoretical analysis to incorporate systems of parallel and interacting investment decisions that occur across units within the firm and among firms.

Second, we find that value-maximising investments in systems of complementary assets require coordination mechanisms that are largely overlooked in recent theoretical literature. In particular, the role of top-level executives extends far beyond Milgrom and Roberts' (1995b, p.13) claim that they "need only identify the relevant complementarity structure in order to recommend a "fruitful" direction for coodinated search" to lower-levels in the hierarchy. At Intel, executives have collaborated with peers in supplier, customer and complementor firms to develop and operationalize a technology roadmap mechanism. We examine how this is used to establish, coordinate and revise expectations, within and between firms, as to when the components of an asset system should be made available and how they should interoperate to enable system-wide innovation.

In contexts where innovation is widely distributed across sub-units and across firms, the benefits of such a coordination mechanism for dynamically adjusting expectations are particularly significant. As we demonstrate for the case of Intel, decisions on accelerating or postponing investments in such as a new microprocessor are embedded in what one executive termed an "ecosystem" (Miller and O'Leary, 2000). Optimal results may be secured only through awareness of proposed shifts in the time-lines and anticipated

outcomes of many other investment decisions, made by such as fabrication process developers within the firm, lithography firms in the supply base, or a set of independent software vendors designing complementary products. The significance of complementarity relations among investments is widely recognized in the literature, and the merits of identifying such relations at intra- and inter-firm levels is also acknowledged. It is important now for researchers to identify and analyse the mechanisms that allow firms to derive the available benefits.

Third, this study identifies issues for investigation in future large-sample surveys and clinical analyses of the capital budgeting process. In particular, it suggests investigating whether there are systematic differences between industries in the effectiveness with which interdependent investments are planned and coordinated across firm boundaries. For instance, anecdotal evidence indicates that firms in the telecommunications industry have found it very difficult to align investments in the components of advanced telephony, with significant negative returns to investment as a consequence (Grove, 2000). If there are such differences across industries, why do they arise? Is it, for instance, due to the absence of appropriate institutional arrangements such as those provided by SEMATECH? Or is it attributable to the lack of a norm such as Moore's Law, through which initial expectations are formed? Or, is it a function of the differing rate and nature of technological progress, such that in one industry (eg microprocessors) innovation is relatively predictable and incremental, and in another (eg biotechnology) it is highly uncertain and fundamental? Further research should focus on whether there are systematic differences across industries with respect to mechanisms for forming, revising and enacting expectations, such that some industries are better able to achieve systemic and inter-firm innovation than others.

As a result of Graham and Harvey's (2001) recent survey, we now have a comprehensive and detailed understanding of the utilisation of particular investment valuation practices and the take-up of real options modelling on the part of large and small firms in a variety of industries. It is important to build upon this information by asking managers whether synergies or complements are addressed formally as part of the capital budgeting process and, if they are, what formal mechanisms are used to do so. Our clinical study suggests the successful use of technology roadmap practices in the computing and microelectronics industries. At Intel, the CEO and other executive officers give priority to investment coordination, with Net Present Value and Net Present Cost calculations taken for granted. This suggests that it is now

appropriate for survey researchers to begin to pose questions relating to other aspects of the capital investment process, and to pay greater attention to issues such as how the relevant unit of investment analysis and appraisal is arrived at. For instance, a roadmap may offer a more robust mechanism for articulating possible responses to the uncertainties of intra- and inter-firm coordination and related uncertainties than that of arbitrarily adjusting the cash-flow forecasts or discount rates of individual investment decisions, an approach which Graham and Harvey (2001) observe is presumed in the existing literature. Systematic investigation of these issues, through survey research in particular, would be of considerable benefit.

Additional clinical studies of the explicit use of formal coordination mechanisms in other industries such as automobile and airplane manufacture would be extremely valuable. It would be of interest to learn whether similar mechanisms to those observed in the microprocessor industry, which allow for the optimising of complementary investments, exist in other industries. It would also be of interest to learn how the coordination of expectations is achieved in other industries. While "Moore's Law" sets out a time-line and a corresponding cost improvement for advances in process technology that is specific to the semiconductor industry, it would be helpful to know whether comparable ways of coordinating expectations with respect to investment decisions exist in other industries.

This clinical study suggests that we need to know much more about how such factors interact with the more traditional aspects of capital budgeting. Investment evaluation and coordination, and the mechanisms through which this is achieved, would seem to be an area where the benefits of applying complementary research methods are particularly strong.

References

Brennan, M., Trigeorgis, L., 2000. Real options. In: Brennan, M., and Trigeorgis, L. (Eds.), Project Flexibility, Agency and Competition. Oxford University Press, Oxford.

Farrell, J., Katz, M., 2000. Innovation, rent extraction and integration in systems markets. Journal of Industrial Economics 48, 413 – 432.

Farrell, J., Saloner, G., 1988. Coordination through committees and markets. Rand Journal of Economics 19, 235 – 252.

Graham, J., Harvey, C., 2002. How do CFOs make capital budgeting and capital structure decisions? Journal of Applied Corporate Finance 15, 1 – 36.

Graham, J., Harvey, C., 2001. The theory and practice of corporate finance: evidence from the field. Journal of Financial Economics 60, 187 – 243.

Grove, A., 2001. Communications strategy and the nation's new agenda. Strategic Management Society Annual International Conference.

Holmstrom, B., Roberts, J., 1998. The boundaries of the firm revisited. Journal of Economic Perspectives 12, 73 – 94.

Jensen, M., 1993. The modern industrial revolution, exit, and the failure of internal control systems. The Journal of Finance 48, 831 – 880.

Jensen, M., Fama, E., Long, J., Ruback, R., Schwert, W., Smith, C., Warner, J., 1989. Editorial: clinical papers and their role in the development of financial economics. Journal of Financial Economics 24, 3 - 6.

Katz, M., 1995. Joint ventures as means of assembling complementary inputs. Group Decision and Negotiation 4, 383 – 400.

Milgrom, P., Roberts, J., 1990. The economics of modern manufacturing: technology, strategy and organization. American Economic Review 80, 511 - 528.

Milgrom, P., Roberts, J., 1992. Economics, Organization and Management. Prentice Hall, Englewood Cliffs.

Milgrom, P., Roberts, J., 1995a. Continuous adjustment and fundamental change in business strategy and organization. In: Siebert, H. (Ed.), Trends in Business Organization. J. C. B. Mohr, Tubingen..

Milgrom, P., Roberts, J., 1995b. Complementarities and fit: strategy, structure and organizational change in manufacturing. Journal of Accounting and Economics 19, 179 – 208.

Miller, P., O'Leary, T., 2000. Value Reporting and the Information Ecosystem. PriceWaterhouse Coopers, London.

Miller, P., O'Leary,T., 1997. Capital budgeting practices and complementarity relations in the transition to modern manufacture: a field-based analysis. Journal of Accounting Research 35, 257 – 271.

Moore, G., 1975. Progress in digital integrated electronics. In: International Electron Devices Meeting Technical Digest. Institution of Electrical and Electronic Engineers, Washington DC.

Trigeorgis, L., 1996. Real Options. MIT Press, Cambridge.

Trigeorgis, L., 1995. Real Options in Capital Investment. Praeger, London.

Tufano, P., 2001. HBS – JFE conference volume: complementary research methods. Journal of Financial Economics 60, 179 – 185.

Wruck, K., Jensen, M., 1994. Science, specific knowledge and total quality management. Journal of Accounting and Economics 18, 247-287.

Appendix A: Intel Corporation, Condensed Financial Statements, 1991 – 2001

Consolidated Statements of Income
(US$ millions)

	2001	2000	1999	1998	1997	1996	1995	1994	1993	1992	1991
Net Revenues	$26,539	$33,726	$29,389	$26,273	$25,070	$20,847	$16,202	$11,521	$8,782	$5,844	$4,779
Cost of sales	13,487	12,650	11,836	12,088	9,945	9,164	7,811	5,576	3,252	2,557	2,316
Gross margin	13,052	21,076	17,553	14,185	15,125	11,683	8,391	5,945	5,530	3,287	2,463
Other operating costs & expenses											
Research & development	3,994	4,006	3,503	2,674	2,347	1,808	1,296	1,111	970	780	618
Marketing, general & admin	4,464	5,089	3,872	3,076	2,891	2,322	1,843	1,447	1,168	1,017	765
Amortization of goodwill	2,338	1,586	411	56	-	-	-	-	-	-	-
Total	10,796	10,681	7,786	5,806	5,238	4,130	3,139	2,558	2,138	1,797	1,383
Operating income	2,256	10,395	9,767	8,379	9,887	7,553	5,252	3,387	3,392	1,490	1,080
Interest income & gains on investment (net of interest expense)	(73)	4,746	1,461	758	772	381	386	216	138	79	115
Income before taxes	2,183	15,141	11,228	9,137	10,659	7,934	5,638	3,603	3,530	1,569	1,195
Provision for taxes	892	4,606	3,914	3,069	3,714	2,777	2,072	1,315	1,235	502	376
Net income	$1,291	$10,535	$7,314	$6,068	$6,945	$5,157	$3,566	$2,288	$2,295	$1,067	$ 819

Consolidated Balance Sheets
(US$ millions)

	2001	2000	1999	1998	1997	1996	1995	1994	1993	1992	1991
Assets											
Property, plant & equipment											
Opening balance (net)	15,013	11,715	11,609	10,666	8,487	7,471	5,367	3,996	2,816	2 ,163	1,658
Additions	7,309	6,674	3,403	4,032	4,501	3,024	3,550	2,441	1,933	1,228	948
	22,322	18,389	15,012	14,698	12,988	10,495	8,917	6,437	4,749	3,391	2,606
Less: depreciation & other adjustments:	4,201	3,376	3,297	3,089	2,322	2,008	1,446	1,070	753	575	443
Property, plant & equipment, net	18,121	15,013	11,715	11,609	10,666	8,487	7,471	5,367	3,996	2,816	2,163
Current assets	17,633	21,150	17,819	13,475	15,867	13,684	8,097	6,167	5,802	4,691	3,604
Marketable strategic equity securities & other investments	1,474	3,712	7,911	5,365	1,839	1,353	1,653	2,127	1,416	496	480
Goodwill & related intangibles	5,127	5,941	4,934	111	-	-	-	-	-	-	-
Other assets	2,040	2,129	1,470	911	508	211	283	155	130	86	45
Total assets	$44,395	$47,945	$43,849	$31,471	$28,880	$23,735	$17,504	$13,816	$11,344	$8,089	$6,292
Liabilities & stockholders equity											
Total current liabilities	6,570	8,650	7,099	5,804	6,020	4,863	3,619	3,024	2,433	1,842	1,228
Long-term debt	1,050	707	955	702	448	728	400	392	426	249	363
Other liabilities	945	1,266	3,260	2,588	3,117	1,272	1,345	1,133	985	553	283
Stockholders' equity	35,830	37,322	32,535	23,377	19,295	16,872	12,140	9,267	7,500	5,445	4,418
Total liabilities & equity	$44,395	$47,945	$43,849	$31,471	$28,880	$23,735	$17,504	$13,816	$11,344	$8,089	$6,292

Appendix B: Effects of Coordinating a Process-Generation Shift with Introduction of a New Product

Panel A	**Panel B**	**Panel C**
Process-generation (x)	Process-generation (x)	Process-generation (x+1)
Product-generation (y)	Product-generation (y+1)	Product-generation (y+1)
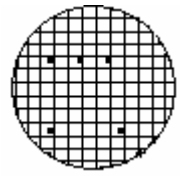	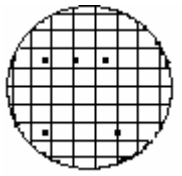	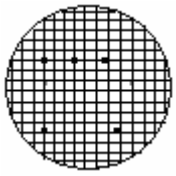

A microprocessor is fabricated by forming electronic elements, such as transistors, on a square of silicon wafer. The elements are connected by layers of metal traces to form a set of integrated circuits. The finished product is a square of silicon embedded with electronic circuitry, termed a die.

Each square on the circles above represents a microprocessor die fabricated on a silicon wafer, and the black dots represent particles that contaminate the wafer during processing, rendering a microprocessor unusable. It is assumed that the number of particles is a function of imperfections in the fabrication process, and independent of the number of die. Each of the three panels shows a total of 5 fatal defects in identical locations.

The shift from panel A to panel B shows the effects of introducing a new microprocessor product without a corresponding change in process-generation. The die-size of product (y+1) in panel B is larger than that of its predecessor, (y) in panel A, because the new microprocessor contains more transistors and circuits to give it added power and functionality. The yield of good-die per wafer is reduced as a consequence: there are fewer die per wafer, and a greater proportion of them are destroyed by the contaminant particles. Fabrication cost per good (or usable) die will rise as a consequence. Also, the clock-speed of product (y+1) may be impaired, because the larger die-size results in electrons travelling longer distances to complete a circuit.

The introduction of the new product (y+1) may be more economic if it is coordinated with a process-generation change, from (x) to (x+1), as represented in the shift from panel B to panel C. The increased transistor density provided by the new process will at least partially offset the increased die-size of the new product, such that the yield of good (or usable) die per wafer and the clock-speed of the device are both increased.