Wireless Data Path for a Mobile, Modular Computer System

**T.D. Kimura, J.R. Gilman,
R.A. Livingston, K. Chan,
and R.D. Chamberlain**

T.D. Kimura, J.R. Gilman, R.A. Livingston, K. Chan, and R.D. Chamberlain, "Wireless Data Path for a Mobile, Modular Computer System," in *Proc. of the 6th Int'l Conf. on Parallel Interconnects*, October 1999, pp. 165-172.

Modubility, LLC
and
Washington University

To appear in IEEE 6th International Conference on Parallel Interconnects (PI'99), October 1999, Anchorage, Alaska, USA

Wireless Data Path for a Mobile, Modular Computer System[1]

T.D. Kimura*†, J.R. Gilman*, R.A. Livingston*, K. Chan*, and R.D. Chamberlain†
*Modubility LLC, St. Louis, Missouri
†Washington University, St. Louis, Missouri

Abstract

We present a comparison of two technologies for use in implementing a wireless data path. The target environment is a mobile, modular computer system that aims to improve the economics and productivity of users that currently use multiple PCs. An inductive power delivery subsystem is described, and both optical and capacitive data delivery subsystems are compared.

1 Introduction

Today's PC platforms do not satisfy individuals' needs for ubiquitous access to information. Desktops are not mobile, laptops are still bulky and too expensive, hand-helds lack computing power, and network computers have not matured. We need a new platform that is mobile and is as affordable and powerful as today's desktop PCs.

In order to develop such a platform, we propose a new scheme for modularizing the PC architecture into a *mobile module* and a *docking module*. The mobile module contains the essential computational resources of the PC, i.e., CPU, memory, and hard disk with OS, but no peripheral devices nor a power supply. The docking module contains the rest of the PC. Our main goal is to develop a mobile, modular PC, which offers PC users economical and productive advantages.

The technical problem facing us is the design of an interface between the mobile module and the docking module. It must provide a reliable, high speed (1 Gb/s) bi-directional data path and power transmission up to 30 W. Furthermore, it must be inexpensive, small sized, and light weight.

In this paper we compare the merits and demerits of two possible solution technologies for the interface design problem. One is optical and the other is capacitor-based. We argue for the capacitor-based technology as a practical solution in the current economic environment, with the optical solution becoming viable as component costs come down.

2 Modularization of the PC

Modularization is a technique for decomposing a complex system into a set of interacting modules, where the functionality of each module is known but its internal structure is not. The merits of modularization have been well understood in the area of hardware component (chip) design, and recently in the area of software engineering, in the form of modular programming [1,2] and object-oriented programming [3,4,5]. For example, Dan Ingals [3] defines *modular programming* as

> A software engineering technique (decomposition method) in which no component in a complex system should depend on the internal details of any other components, i.e., a black-box approach for software engineering.

Webster's dictionary defines *modular* as 'being constructed with standardized units or dimensions for flexibility and variety in use.' Modularization facilitates the reusability of modules and the flexibility of system construction. It enhances the system reliability by reducing the inter-dependency among the modules and by preventing error propagation. In short, modularization promotes the

[1] Information contained in this paper is the subject of and protected by one or more patents, patent applications and/or other forms of intellectual property rights. No license, actual or implied, is granted under those rights.

system manageability and reduces the system cost by sharing standard modules with others.

A principle for a successful modularization (decomposition) consists of two parts:

1) The inter-module interface must be uniform and standardized so that the modules with the same functionality may be guaranteed to be inter-changeable.
2) Tightly coupled components are clustered into a module so that the inter-module coupling (traffic) may be minimized.

At the user level, a typical desktop PC is modularized into the CPU chip, memory chips, a motherboard, a hard disk drive, and various peripheral units. The interface between the CPU and the motherboard, a socket, is standardized by chip manufacturers. Similarly, for the memory chips, standard interfaces are used. The interface between the peripheral units and the motherboard is standardized industry-wide as PCI, ISA and/or USB.

At the system level, the motherboard is decomposed into a high-speed *system bus* which facilitates the interaction between CPU and memory, and a *peripheral bus* which connects the peripheral units to the system bus through a bus bridge. In many of today's desktop PCs, the clock speed of the system bus is 100 MHz, while it is 33 MHz for the peripheral bus (PCI).

3 Mobile, Modular PC

We propose to introduce another modularization scheme at the user level, applying the principle of modularization to a decomposition of the motherboard. In this scheme, the PC architecture is bisected at the bus bridge into a *mobile module* and a *docking module*. The bus bridge itself resides at the boundary between the mobile module and the docking module. The mobile module contains a CPU, memory, and a hard disk as well as a system control chipset and a clock, but nothing else. The docking module contains the rest of the PC including the power supply. We have constructed a prototype of a mobile module and a docking module using off-the-shelf components. The photograph below (Figure 1) shows the mobile module circuit board (from Ampro [6]) on the left and hard disk (from IBM) on the right.

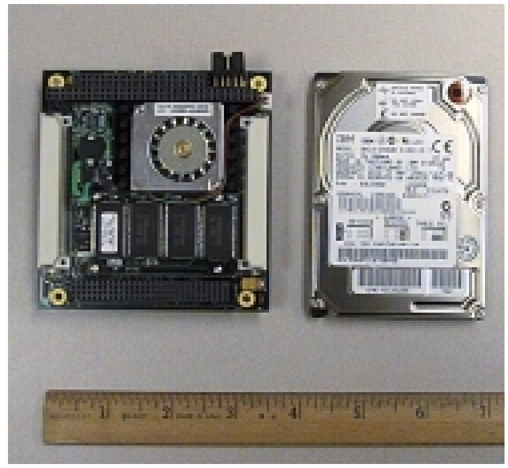

Figure 1: Mobile Module (Prototype)

The figure below (next page, Figure 2) gives the block diagram of the prototype. The mobile module contains a 133 MHz Pentium processor, 64 MB of memory, a 2.5" hard disk with 4 GB capacity, and a Northbridge chipset for the PCI bus. When assembled, it measures 4" x 5" x 1.5" and weighs approximately 10 oz. The docking module contains the display subsystem, USB controller (with attached keyboard and mouse), CD-ROM drive, network interface, power supply, etc.

This modularization is founded upon the following observations: (1) CPU, memory, and hard disk (containing an OS) are the essence of a computing environment. (2) They are the most tightly-coupled components in the PC architecture and their performance is mutually dependent. (3) They are the most expensive parts of the PC. (4) Their total weight is lower than any other peripheral devices. Thus, the primary purpose of this scheme is to provide PC users with the portability and mobility of the essence of the PC, i.e., OS, application software and user documents. A power supply is excluded from the mobile module for the same mobility reason.

There may be a variety of docking modules, from a desktop unit to a palmtop docking unit. Every docking module should share the same interface so

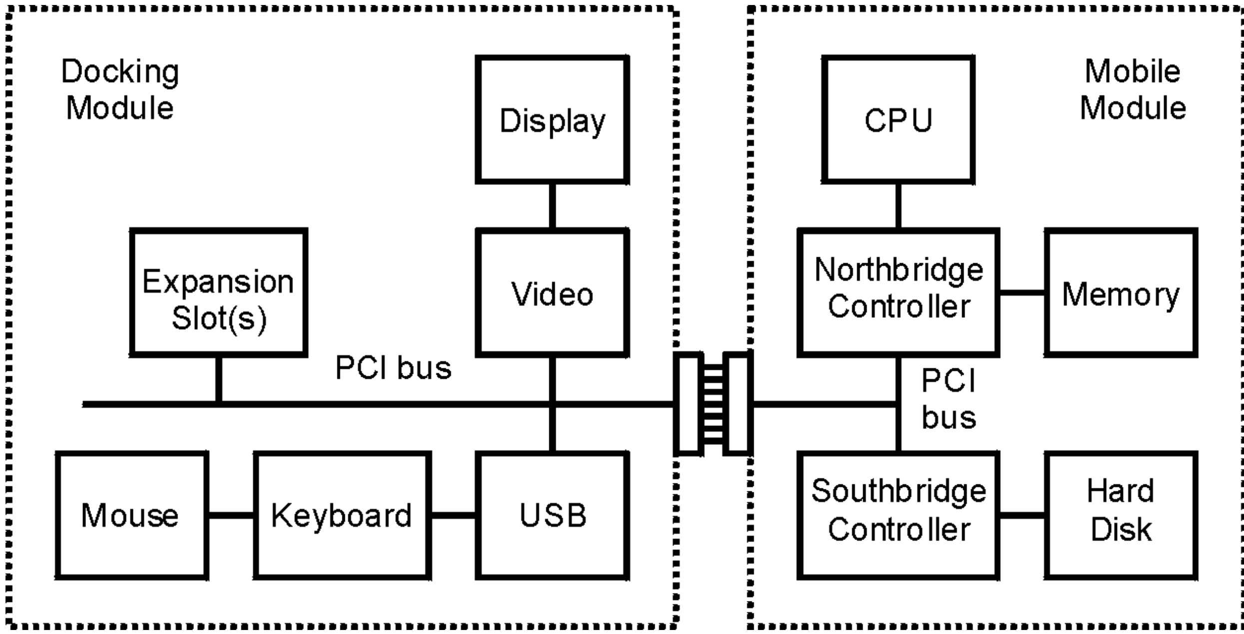

Figure 2: Block Diagram of Prototype

that a mobile module can be connected to any type of docking module. The end user carries a mobile module in his/her pocket and, whenever PC access becomes necessary, the user docks it to a docking module that is either a stationary desktop unit, or a mobile unit equipped with an LCD display panel and a battery. In a sample usage on a college campus, students own a mobile module, and colleges provide an infrastructure of desktop docking units, in the library, classrooms, and dormitory rooms. The average frequency of docking/undocking is expected to be high, 10 – 20 times a day on campus.

This modularization is motivated by the following merits to the end user who needs ubiquitous access to multiple PCs:

- *Reduced Cost*: The total computing cost can be reduced by eliminating duplication of otherwise necessary hardware (multiple CPUs, memory, hard disks, etc.) and software (both OS and applications).
- *Easy File Management*: The file synchronization problem is completely eliminated.
- *Consistent User Interface*: The user need not learn different user interfaces (desktop environment and directory structure) at different PCs. The user carries a single user interface with him/her.

4 Interface Requirements

In the current prototype, the mobile module and the docking module are connected by a mechanical 120-pin connector. It is clear that such a mechanical link will become unreliable after frequent docking/undocking operations, not to mention the awkwardness of aligning multiple pins during each docking operation.

What is needed is an interface technology that is reliable, easy to use, small sized, light weight, and economical. One possibility is a wireless (contactless) interface. The figure below (next page, Figure 3) illustrates the use of a wireless interface between the mobile module and the docking module. Note that we consider a wireless link not for its mobility but for the increased reliability without mechanical wear and tear. Therefore the maximum distance between the two modules can be limited (e.g., 1 to 2 mm). We will investigate the technological issues associated with a wireless interface in the next section.

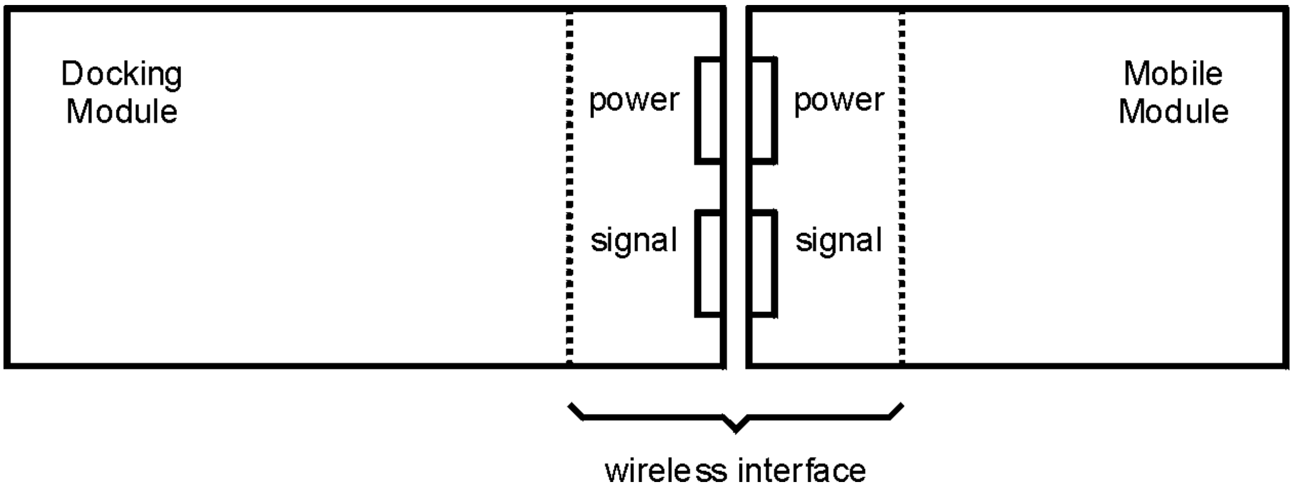

Figure 3: Wireless Interface

We list below the requirements specification for such an interface:

- Reliability: The interface has to accommodate a high frequency of docking/undocking operations without wear and tear.
- Ease of Use: The user should be able to dock/undock in a second without any alignment problem.
- Signal Bandwidth: The interface must provide a bi-directional data path of at least 132 MB/s. This is based upon current peripheral bus speeds (e.g., PCI).
- Power transmission: The interface must provide power transmission of at least 30 W. A mobile module with a high speed CPU (e.g. Pentium II) will require this level of power.
- Cost: Reduction of total computing cost being the main motivation of our modularization, the incremental cost of the interface must not exceed 10% of the cost of a mobile module.
- Physical Dimension: In order to maintain the mobility/portability of the mobile module, we need to minimize the physical dimension of the mobile module and the interface itself.

5 Interface Technology Feasibility

The above sections provide the motivation and requirements for a wireless interconnect that provides both power and data transmission between separately packaged portions of a computer system. This section addresses the technological feasibility of such an interconnect. We first describe an inductive power delivery subsystem, and then describe two mechanisms for data delivery: an optical link and a capacitive link. Prototype designs of each of these subsystems have been constructed and their operation verified.

5.1 Power

The power delivery subsystem uses a split-core transformer to deliver power across the inter-package gap. This is illustrated as follows:

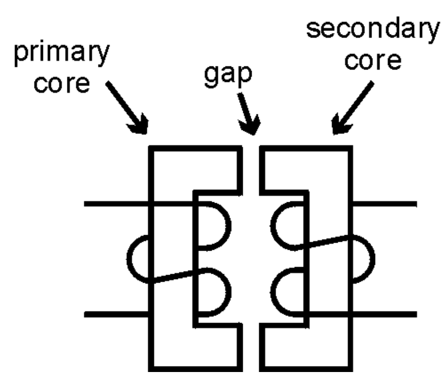

Figure 4: Wireless Power Transmission

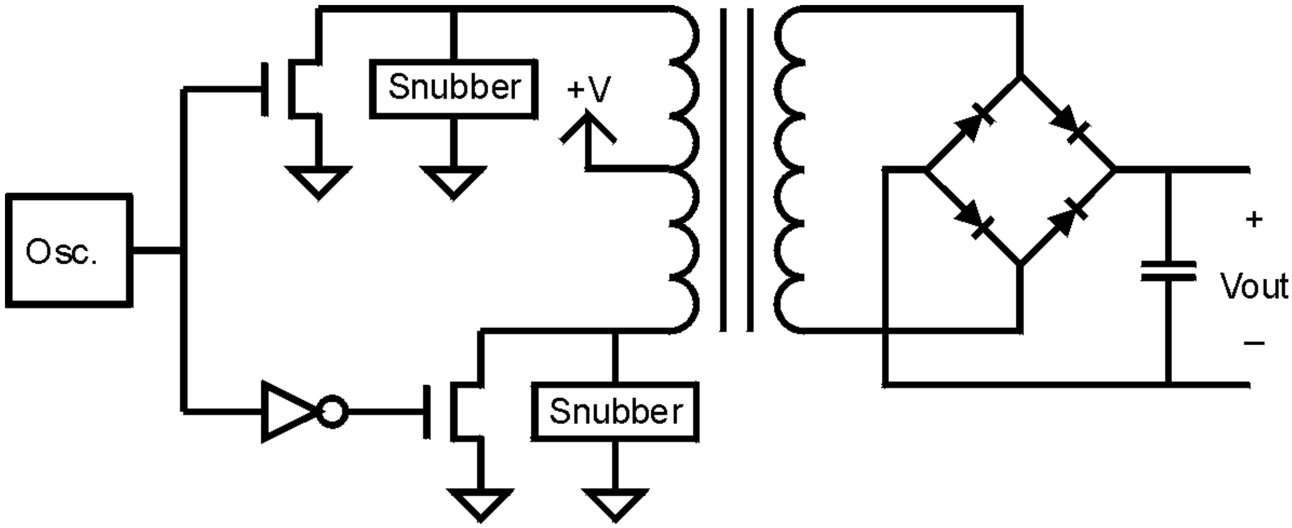

Figure 5: Wireless Power Transmission (Schematic)

The primary core is positioned in the docking module, and the secondary core is part of the mobile module. The gap between the cores contains the enclosure side walls (moldable plastic). An idealized schematic is shown above in Figure 5.

The snubber circuits are essential in a split-transformer design to dissipate the magnetization currents present in the transformer primary. These currents are significantly higher than in traditional power supply design, since the effective inductance decreases as the gap distance increases. The circuit is designed for operation up to 35 W delivered to the load with a 20 mil gap and has been successfully tested up to 20 W delivered to the load. The photograph on the left (Figure 6) shows the transformer primary and secondary core and windings.

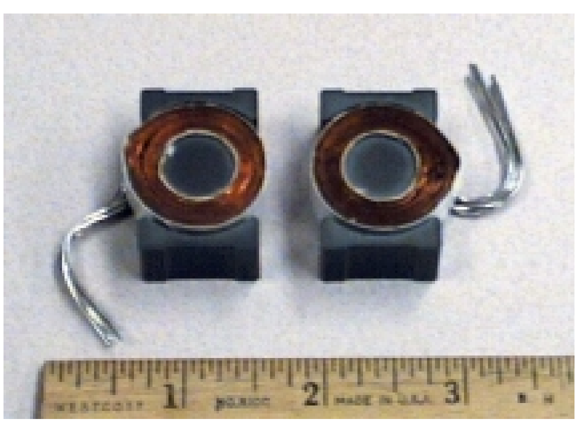

Figure 6: Split Transformer (Prototype)

5.2 Optical Link

The first data delivery mechanism to be described is an optical link. The basic topology for a bi-directional link is illustrated below in Figure 7:

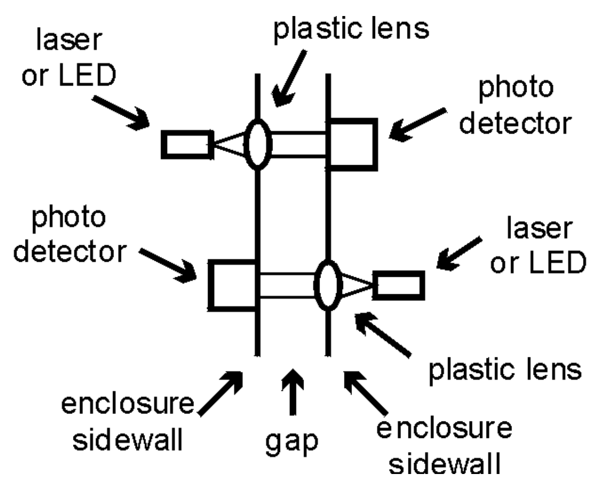

Figure 7: Topology for Bi-directional Link

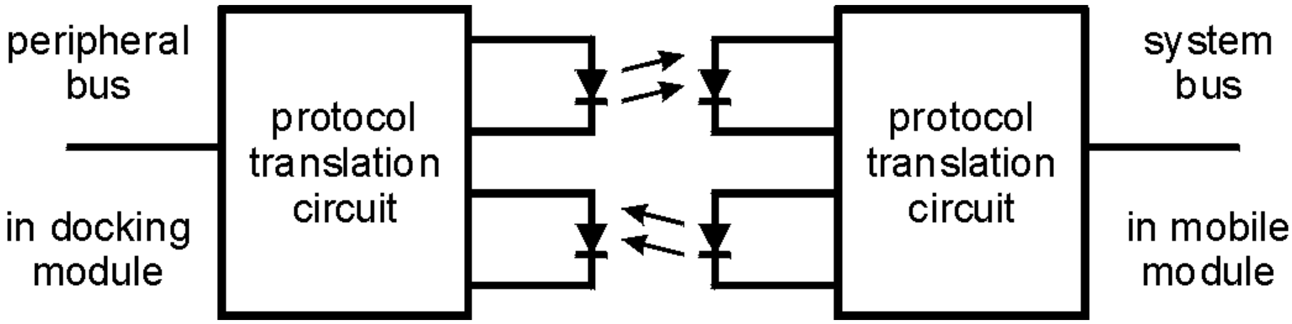

Figure 8: Use of Optical Link

The left-to-right data path is shown at the top of the figure, and the right-to-left data path is shown at the bottom. Each link consists of a light source (laser diode or LED), a small plastic lens to collimate the light from the source, the gap between enclosures, and a photo detector. Clearly, the enclosure material must either be transparent to the wavelength used or openings must be provided. The optical design is fairly straightforward, especially since the power budget can be extremely generous. With gap dimensions on the order of a millimeter or less, only a small fraction of the light need be incident on the detector for reliable operation. We have prototyped the above design using 100 Mb/s Ethernet signals as the data source and sink.

The figure above (Figure 8) illustrates the use of the optical link in the context of the modularized computer architecture.

At the left is the peripheral bus in the docking module, and at the right is the system bus in the mobile module. The circuit as a whole bridges the two (potentially distinct) busses. The optical link bandwidth must be consistent with the peripheral bus data rate (e.g., 132 MB/s for current PCI bus). If it is bit serial, the protocol translation function must include data serialization for the transmitter and serial to parallel conversion at the receiver.

There are a number of difficulties associated with the optical data link. First, optical components capable of Gb/s data rates are available, but their cost is significant. Second, parallel to serial and serial to parallel data conversion adds latency to the protocol translation functionality. On the plus side, the generous optical power budget allows for a significant degree of positional tolerance.

5.3 Capacitive Link

An alternative to the optical link described above is the use of capacitance to deliver the data signals. Here, a conductive plate is attached to (or molded into) each of the enclosure sidewalls,
and when the mobile module is docked, the two plates form a capacitor. The dielectric of the capacitor is the enclosure material (typically ε_r is about 2). The link is illustrated below:

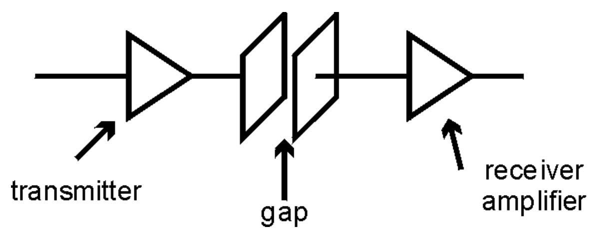

Figure 9: Capacitive Link

Using the standard equation for a parallel plate capacitor [7]:

$$C = \varepsilon_0 \varepsilon_r \frac{A}{d}$$

and assuming 5 mm by 5 mm plates and a 0.5 mm gap, this yields approximately 0.9 pF of capacitance. This implies only a small signal is available at the input to the receiver amplifier.

In the prototype implementation, the receiver amplifier has sufficient hysteresis that it retains the last state transition delivered across the capacitor.

The ability to integrate the transmitters and receivers, combined with the low cost of the capacitors themselves, implies that a parallel data path is entirely appropriate here. The primary limitation in the degree of parallelism is the total area consumed by the capacitors. For a 32-bit data path and the 25 mm^2 capacitors described above, a total area of 1600 mm^2 is required (assuming one half of the area is consumed by the capacitors).

The integration of the capacitive link into the computer architecture is shown in the figure that follows. In order to conserve area, each capacitor is used in a bi-directional fashion, with both a transmitter and receiver circuit on each side of the link. To enable this to work, the drivers must be capable of disabling their output (i.e., a high-impedance state). The result is that the capacitive link acts very much like a traditional bus, bi-directional in nature with drivers that can be turned on or off depending on whether or not they are needed. As in the optical link, the protocol translation circuits are responsible for interfacing to the appropriate bus on either end of the link. An additional requirement for the capacitive link is the need to encode the link data such that there is no DC signaling present.

The disadvantage of the capacitive link, relative to the optical link, is the fact that the gap must be held to a minimum so as to not overly diminish the capacitance values. In the optical link, there isn't as strong a need to maintain tight separation between the mobile module and the docking module. The main advantage to the capacitive link is economy.

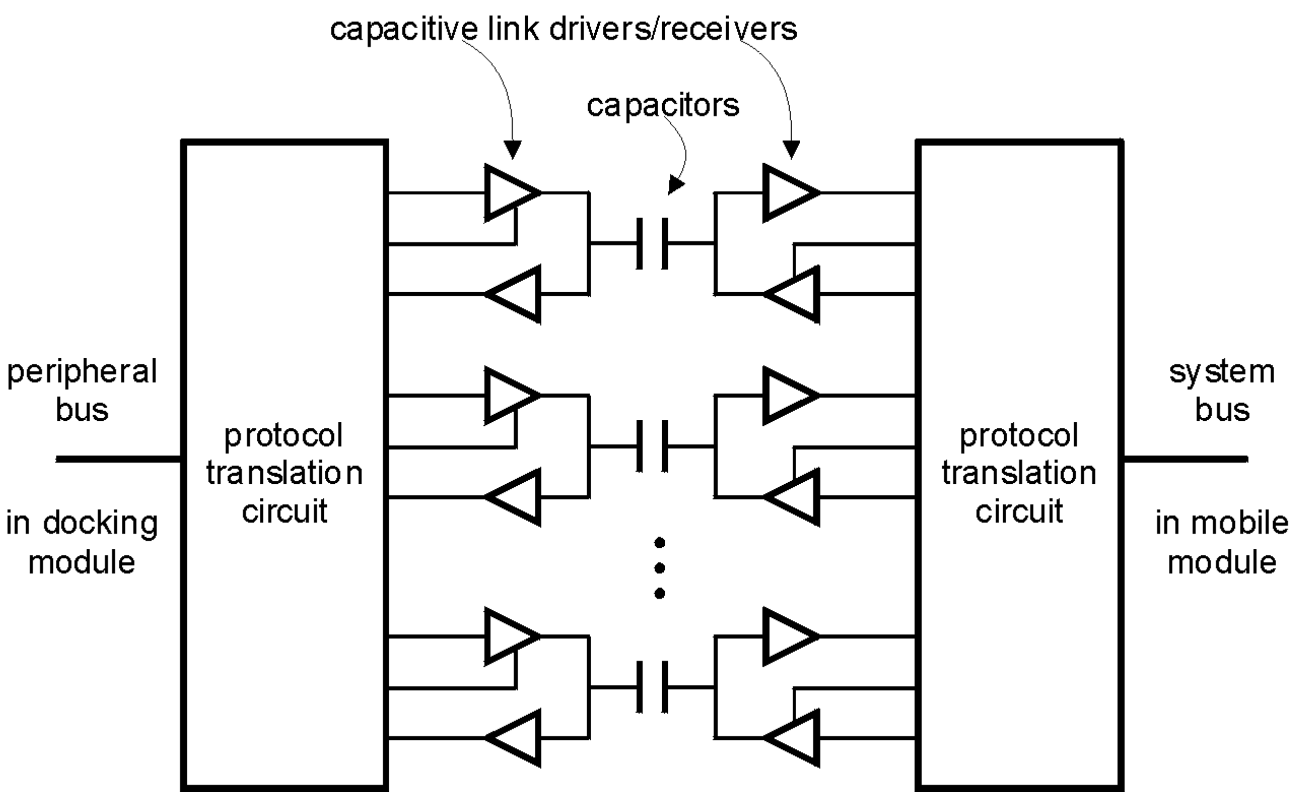

Figure 10: Capacitive Link in Computer Architecture

6 Conclusions

This paper has described an approach to modularizing the PC architecture to improve its mobility, economy, upgradeability, and flexibility. As part of this modularization, we have investigated several technical issues associated with the inter-module interface. The high-level results of this investigation are summarized in the table below (Table 1).

In the current economic environment, the high cost of the optical link essentially precludes its adoption, making the capacitive link the most attractive.

Generally, as the cost of optical components (specifically diode lasers) decreases, the attractiveness of the optical link increases. With VCSELs as the light source, one can consider a parallel interface, eliminating the need for parallel-to-serial and serial-to-parallel data conversion.

7 References

[1] Niklaus Wirth. *Programming in MODULA-2*, Springer-Verlag, 1982.
[2] Martin Reiser and Niklaus Wirth. *Programming in OBERON, Steps Beyond Pascal and Modula*, Addison Wesley, 1992.
[3] Daniel Ingals. "The Smalltalk-76 Programming System: Design and Implementation," *Proceedings of the ACM Symposium on Principles of Programming Languages*, January 1978.
[4] Bjarne Stroustrup. *The C++ Programming Language*, Addison-Wesley, 1991.
[5] James Gosling, Bill Joy, and Guy Steele. *The Java Language Specification*, Addison-Wesley, 1996.
[6] Ampro Computers, Inc., CoreModule/P5I Manual, San Jose, CA, 1998.
[7] Richard C. Dorf, ed., *The Electrical Engineering Handbook*, CRC Press, 1993.

Requirements	Wired Link	Optical Link	Capacitive Link
Reliability	Low	High	High
Bandwidth	High	High	High
Power delivery	Straight-forward	Requires split-core transformer	Requires split-core transformer
Positioning tolerance	High	High	Low
Cost	Low	High	Low

Table 1: Comparison Results