Implementation Documentation for the MPC79 Multi-University Multiproject Chip-Set

A Collection of Information and Instructions conveyed to the Participating Designers, along with their Packaged Chips.

Compiled by

Lynn Conway, Alan Bell, Martin Newell, Richard Lyon, and Richard Pasco

LSI Systems Area, Systems Science Laboratory, Xerox Palo Alto Research Center.

1 January 1980

During the fall of 1979, the LSI Systems Area of Xerox PARC/SSL conducted the demonstration-operation of a prototype system for remote-entry, fast-turnaround implementation of VLSI designs. The user community for this demonstration was composed primarily of EE/CS students taking courses in VLSI system design at many major universities throughout the United States.

A total of 82 VLSI system design projects, created by a total of 124 participating designers, are included in the MPC79 chip-set. Implementation began at the design cutoff-time of 5:00 pm PST, 4 December 1979. Packaged chips, custom wire-bonded for each project, are being distributed to all participants on 2 January 1980. Thus the turnaround time for the MPC79 chip-set was 29 days.

This document provides basic information that the participating designers may use (i) to determine the package-codes for their projects, so as to identify which of the boxed, packaged chips contains their project, (ii) to determine their project pinouts and prepare their chips for testing, (iii) to estimate the maximum clock rates for their projects, (iii) to study the details of the starting frame, and (iv) to learn more about the other participants and the general results of this happening.

Some plans for further documentation and publication concerning MPC79 are also described. Included is a request for short reports (an example short report is appended) from those who'd like to see their projects featured in an upcoming *MPC79 Proceedings*. Many of the projects included in MPC79 have very imaginative, innovative architectures and novel design features. We strongly encourage the designers of such projects to seek publication of their work in the technical journals.

List of Sections:

[Filed on: [MAXC]<Conway>MPC79imp.doc, January 1, 1980 7:45 PM]

1. Introduction and Overview

This document provides a collection of information and instructions for use by the VLSI system designers participating in the MPC79 multi-university multiproject chip-set.

This introductory section provides background information about MPC79, about the context in which it occurred, and some reasons why the effort was undertaken. Summaries are provided of the participating universities, of the organizations involved in the implementation effort, and of the overall schedule of MPC79 events. This background material may help provide the general reader with a visualization of what MPC79 was all about.

Section 2 contains important basic information that the designers can use to identify their packaged project chip and to prepare the chips for testing. Section 3 contains project locator maps and a complete list of all projects and designers. Sections 4 and 5 provide information about the starting frame and electrical test results. Section 6 presents a complete set of unmarked wire-bonding maps. These maps can provide the designers and the general reader with a visualization of the scope of the MPC79 projects and the overall effort. Finally, Section 7 discusses our plans for packaging and distributing more chips, for taking and distributing chip photographs, and for writing additional reports concerning MPC79.

The Background and Context of MPC79

Since 1976, the LSI Systems Area of the Systems Science Laboratory (SSL) at Xerox Palo Alto Research Center (PARC) has conducted research in the architecture and design of integrated systems. One focus of that research has been the exploratory development, in collaboration with the Caltech Computer Science Department, of new design methodologies which simplify integrated system design so that this capability can be more quickly acquired and more widely practiced by system designers than was possible in the past.

Through this research, new design techniques have evolved and have then been debugged in actual practice in the university classroom, and during in-house short courses. This has been a process of discovery and iteration, moved forward by teaching design techniques to university EE/CS students and to practicing EE's and computer scientists, and then having the students undertake VLSI design projects as part of their course experience. Groups of these VLSI projects, and thus the design methodology itself, have been subjected to the "acid test" of actual implementation and testing.

Our teaching and implementation activities have thus been an integral part of our research methodology. There has been success in this effort, leading to the publication of the textbook "Introduction to VLSI Systems", by C. Mead and L. Conway, Addison-Wesley, 1980, describing the new design techniques. A number of major universities, including MIT, CMU, Stanford, Caltech, and U.C. Berkeley, now offer standard courses in VLSI system design based on this textbook.

By early 1979 it had become clear that a demand for fast-turnaround VLSI implementation would develop that was far beyond our ability to supply using our early ad-hoc methods. We then conceived the idea of developing what amounts to an "operating system" to interface a large number of remote users with centralized maskmaking, wafer fabrication, and packaging facilities, making automated responses to electronic messages from users, juggling the many constraints involved, and handling the vast logistics required to coordinate the overall implementation effort.

We were convinced that such a system could significantly reduce implementation turnaround-time, could reduce by about two orders of magnitude the cost per design-iteration compared to that normally incurred in the industry, and could bring access to the fabrication process to a much larger technical community than in the past. We conceived of evolving such a system using the same method used to develop the design methodology: By design, use, and iterative debugging in the universities, with the system used to provide an implementation service to university designers.

The architecture of a prototype implementation system was undertaken and completed during the summer of 1979 by Alan Bell and Martin Newell. They completed the design and coding of the system during the early fall of '79, in time for use to support the fall 1979 VLSI design courses.

The MPC79 effort was then mounted, for the purposes of (i) continuing our support of university VLSI design courses by providing the implementation of student-project designs, (ii) demonstrating the feasibility and general capabilities of such VLSI implementation systems to a wide technical community, and (iii) refining our ideas concerning the architecture, design, and operation of such systems, by running a major operational test of a prototype system.

It is unusual and perhaps somewhat controversial for us to mount such vast experiments in a large university community. There are great risks involved: risks of failure and the risks associated with presenting undebugged technology and techniques to a large group of students. However, we've found the universities eager to collaborate and to run these risks with us. It is exciting and we believe it is appropriate for university students to be at the forefront, sharing in the great adventure of developing and applying new knowledge.

The Participating Universities and Designers

The user community for this demonstration was composed primarily of EE/CS students taking the courses in VLSI system design at major universities throughout the United States, along with a number of university faculty and research staff members undertaking major VLSI system designs.

The MPC79 chip set contains a total of 82 VLSI system design projects from a total of 124 participating designers. Designs were included from VLSI design courses at M.I.T., Caltech, Stanford University, Univ. of Illinois, and Univ. of Rochester. MPC79 also includes a number of designs by faculty and research staff members at M.I.T., CMU, Stanford, U.C. Berkeley, Univ. of Washington, Yale University, Univ. of Bristol (England), and Univ. of Colorado (Colorado Springs).

The project-oriented university VLSI design courses were run on a very tight schedule: The courses began in mid-September. By early November, students had learned enough about the basics of VLSI design to originate a project concept and begin design work. The design cutoff date was 4 December 1979. Thus, students learned how to design in about 6-7 weeks, and then completed a project in the remaining 6-7 weeks of the course. Most large research designs in MPC79 took somewhat longer to create, and some of these are iterations of earlier designs by the researchers.

This has resulted in quite a wide range of interesting designs, as you can visualize by looking at the maps in Section 6. They range from small individual student projects, on through to large efforts by groups of designers. Some are very large, ambitious, team projects by university researchers, using sophisticated design aids of their own creation.

MPC79 was run on a schedule to deadlines, in order for fast-turnaround maskmaking and wafer fabrication to be scheduled. We did not censor designs: If a design met certain prearranged requirements for space allocation, consisted of syntactically correct code, and was submitted by the coordinators as "finished" before the deadline, then it was included. In a few cases, designers took on a bit too much for the available time (if you look closely at the Sect. 6 maps, you'll find cases where last-minute panic changes led to glorious disasters). But that was all part of the happening!

Organizations involved in the implementation effort

Several other organizations collaborated closely with us in conducting this demonstration: Data communications (electronic messages, design file transfers) were supported using the ARPANET; maskmaking was done by Micro Mask, Inc., using an electron-beam maskmaking system; and wafer fabrication was done by Hewlett-Packard's Integrated Circuit Processing Laboratory. Further information on this university-industry-goverment collaborative effort will be given in an MPC79 Proceedings and in several technical reports, to be published in the spring of 1980 (see Section 7.).

For an overview of the flow of information and artifacts through these various organizations, see the MPC79 Flowchart at the end of this section (that same flowchart, along with a page of text describing MPC79, is printed on the "document chips", one of the chip types in each wafer set!).

MPC79 implementation schedule

Following the design cutoff time of 5:00 pm, Tuesday December 4, we started final processing of the implementation requests. At 10:00 am the next morning we took the merged mask-specification data to Micro Mask. Delivery of the masks was pipelined with the early fabrication steps, with the first mask of both mask sets being delivered at 5:00 pm, Thursday Dec. 6. By then, HP-ICPL had already started the processing and were ready for the first masks. Processing continued normally except for one major contingency: the failure of a poly-deposition system (causing a delay of about 8 days for repairs, and leading to an additional delay in wafer processing due to the Christmas holidays). At 10:00 am on Dec. 28, the B wafers were ready, and on Monday Dec. 31 the A wafers were ready. Packaged chips, with custom wire-bonding maps, were shipped out to all participants on 2 January, 1980. Thus the implementation turnaround-time for MPC79 was 29 days.

Most of this document was itself produced automatically by the MPC implementation system, directly from the design-file data-base and archived records of the MPC79 die-layout-planning/design-merging process. This method of rapid document-creation enabled us to return this material in a timely way to the designers, right along with their packaged chips.

Looking Ahead

Many of the MPC79 designers are already looking ahead, making plans for iterating their present designs, and dreaming up even more ambitious design projects. Many other university students will want access to VLSI implementation, so that they too can have the experience of learning to design in a state-of-the-art technology by actually *doing it.* The demand for such services may build

rapidly, once folks know that it is feasible and can visualize how small is the expenditure of resources per chip-set when compared to the value of the result to the community of designers. Thus, those designers who'd like to iterate their present designs, or take on a larger design, should take heart! There may be several MPC's in 1980! You designers can help to make sure this happens by letting other folks know what you've done, and how it was done.

It is our sincere hope that MPC79 has provided a sufficient demonstration of the feasibility and practicality of remote-entry, fast-turnaround VLSI implementation, that it will lead to the funding and operation of a regular, scheduled VLSI implementation service for university students and researchers. We believe such a service will achieve an enormous return to the country on its investment, by greatly leveraging the human resources to be applied in the exploration of integrated system architecture and design.

Acknowledgements

We wish to express our gratitude to all the folks who pulled together with us to make MPC79 happen, including our friends at Defense ARPA, at Micro Mask, and at Hewlett-Packard/ICPL. We all owe a great deal to the dedication and efforts of the instructors and lab coordinators in the universities, who, working under great pressure and on a tight schedule, have done such a fantastic job with their design courses and project-labs this fall. We especially want to thank all the student designers for their enthusiastic response to the courses, and for their magnificent efforts and accomplishments on their VLSI system design projects. You have done well! Organizing and carrying out the implementation of these many imaginative VLSI design projects has been a very exciting and rewarding experience for us here at Xerox PARC/SSL.

MPC79 Flowchart:

DS 12; 9 PlaCell;
(5 Items.);
L NM; B L 4000 W 1000 C 2000, -750;
. . .
L NP; B L 500 W 4000 C 2500. -2000;
DF;

TO: MPC79@PARC-MAXC
FROM: REB@MIT-XX
SUBJECT: IMPLEMENT PROJ.CIF

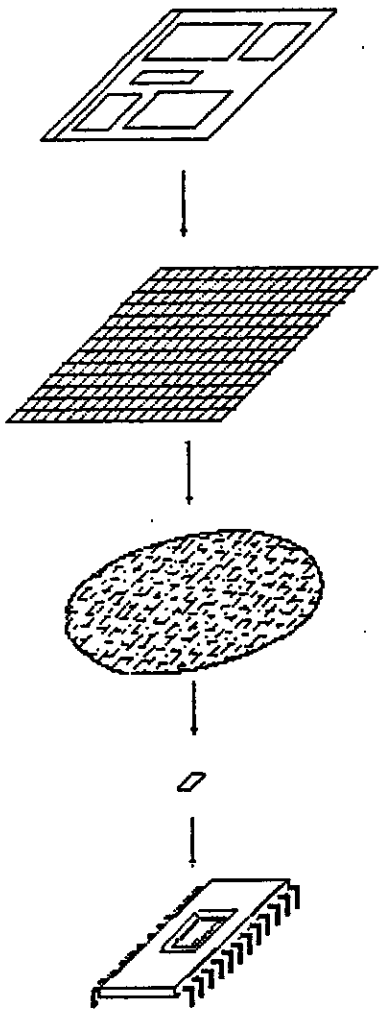

USER COMMUNITY

~ 100 Designers at:
MIT, Caltech, Carnegie-Mellon Univ., Stanford,
Univ. of Illinois, U.C.Berkeley, Univ. of Wash., - - -
(using AIDS/LAP/ICARUS/etc..)

(MSGS, Design Files)

Project lab coordinators at each school use local electronic mail and file transfer facilities to interact with the designers and use the ARPANET to interact with MPC79

DATA COMM. FACILITY

(MSGS, CIF2.0 Design Files)

ARPANET
(MSG, FTP, TELNET)

(MSGS, CIF2.0 Design Files)

IMPL.FACILITY

INFO.MGMT SYSTEM: XEROX PARC/SSL
Checking, Planning, Merging of Designs into Starting Frames
Meeting of Constraints, Coordination , Logistics

(Constraints, logistics, control info.)

(Design files, merged into Starting Frames)

"The Foundry"

MASKMAKING: MICRO MASK, INC.
Data format: MEBES; ETEC electron-beam system

(MASKS)

WAFER FABRICATION: H-P/ICPL
NMOS Silicon Gate
LAMBDA = 2.5 MICRONS

(WAFERS)

PACKAGING

(Bonding Maps) (Packaged Chips) (Elect. Params) (Plots)

Packaged Chips, custom wire-bonded per project, along with plots, wire-bonding maps, and results of electrical testing, to send back to the designers for functional testing.

2. Some Basic Instructions for the Designers:

This section contains practical information that designers may use (i) to figure out which of the boxed, packaged chips is theirs, (ii) to become aware of cautions to observe when handling their chips, (iii) to determine the pinout of their project, and (iv) to prepare their project for testing.

How to find the packaged chip containing your project

Each MPC79 project has a unique Wafer-type/Die-type/Project-number code, as a function of where that project is physically located on the wafers (for example, Project number 5, on Die-type K, from Wafer-set B, is labelled "BK-5"). Each packaged chip is labelled with the Wafer-type/Die-type/Project-number code of the project within the chip that is bonded to the package leads. The plastic box holding the packaged chip, and the chip package itself are both labelled with this code.

You can find the code for your project in the lists in Section 3 of this document, where these codes are associated with the MPC79 Project IDs (ex: project "BK-5" has the ProjectID "Clark2SU", and belongs to Jim Clark of Stanford University). So, all you have to do to find your packaged chip is look up the code in the lists in Section 3, and then locate the box marked with that code. You should also obtain the custom wire-bonding map marked with your code.

Do be careful when handling the chips

Use caution when handling the boxed chips. Avoid opening and closing the boxes unnecessarily, in order to avoid accumulation of dust and dirt on the chip within. Dust and dirt, although not necessarily damaging to the chips, mars their appearance under microscopes and in photographs. (Note that the chips can be studied under low-power stereo microscopes by looking right through the plastic box cover, with the cover closed). If you open and close the boxes often, even in a dust free environment, an accumulation of tiny white plastic blobs will build-up inside the box and on the chip (these come from the box catch). These won't hurt anything, but they look rather wierd under the microscope. These blobs, and dust particles, can be cleared off the chip's surface (to prepare for photos, etc.) by carefully blowing air over the surface with a plastic squeeze bottle.

Be *especially* careful when handling your project chip outside of the box, *so as to avoid breaking or shorting the bonding wires.* When you are ready to test your project, and especially if you intend to place it in operational use, you should fashion some sort of removable cover to tape over the package cavity. *However, if you do cover the cavity, be sure not to press down on the bonding wires.* In some cases the bonding wires protrude above the plane of the package surface. Provision of a workable cavity cover will require some ingenuity in such cases.

Until and unless you prepare your project for testing, it is best to keep it in its plastic box. Note that the piece of cardboard wrapped around the package does two things: It provides a tight fit for the package, keeping it from sliding around inside the plastic box, and it also prevents the bonding

wires from being smashed by the box cover. So, install the cardboard wrapper on the package before inserting the package into its plastic box.

Preparing your chip for testing

There are a few simple steps required to prepare your project for testing. First, you will have to figure out your project's package-pinout, in order to know how to interface it to your test equipment. A custom wire-bonding map for each project has been returned with the packaged chips. Find the one with your project code. That map will show how your project has been wire-bonded to the package leads. The map will enable you to deduce the association of project-pads and package pin-numbers. To locate particular package pins, use the following pin locator diagram:

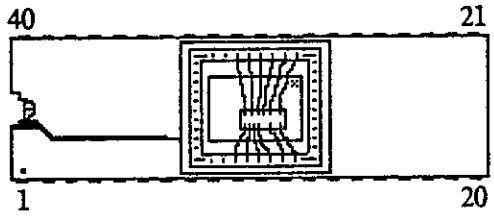

The diagram shows the top view of a ceramic, 40-pin, sidebrazed package having a cavity size of 340 x 340 mils, containing an MPC79-sized project chip (of die-type "X"). One project is bonded-up to the package leads, with wires bonded to leads 9-14, and 27-32. The sketch indicates the positions of the external package pins, with Pin 1 in the lower-left marked by a dot on the package.

The package has some interesting features: Pin 1 is connected to the metal base of the cavity in which the chip is glued (using conductive epoxy). Thus Pin 1 can be used to ground the substrate, or to set it to some particular negative bias. However, it is OK to leave Pin 1 unconnected (if you don't hook-up Pin 1 to some external connection, the substrate will float down to some modest negative voltage during operation of the chip). If it is necessary to use all 40 pins, then the Pin 1 connection to the substrate can be broken by breaking off the "breakaway notch" at the left end of the package. Two of the MPC79 projects use all 40 pins, and require this notch to be broken. One MPC79 project, AF-1, has 48 bonding pads, and will be mounted in a different type of package.

Now, you'll note that each row of package pins are all hooked together by a metal strip. These strips are an artifact of the package manufacturing process. In order to prepare the package for insertion into test fixtures, you'll have to cut these metal strips off the pins. This is easily done using a set of sturdy, sharp scissors, or a pair of sharp metal-nippers. To keep the package from sliding around in its box following removal of these strips, place some plastic foam in the box (and unless a cover has been placed over the package cavity, continue to use the cardboard wrapper).

If you plan to leave your chip installed in test circuitry for any period of time, it's a good idea to tape some sort of cover over the cavity (but be careful not to press down on the bonding wires!). This will protect the chip and the bonding wires. The cover may also be necessary in certain cases to reduce the light reaching the chip during in-circuit operation (see Mead & Conway, p. 134).

The starting frame and electrical test data later in this document will help you estimate the expected maximum clock rate of your project. Your chip is now ready for functional testing!

Selecting chips for further packaging

In addition to the boxed, packaged chips, we've shipped some unpackaged chips to each school. More will be returned in the coming months. Lab coordinators or designers having access to packaging facilites may wish to package some of these chips. To select chips for packaging a particular project, first find some chips of the correct die-type code (the die-type code is indicated by a large letter in the upper-right corner of each chip). Then make an optical inspection under a high-power microscope of the desired project on each of the chips. You may want to make this inspection after the chips are mounted in packages, to detect any damage caused during mounting. If no obvious defects are seen in that project on a chip, then that chip is suitable for the bonding of the project. If there is an obvious defect in that project on a chip, that chip is not discarded but is instead returned to the collection of unpackaged chips (perhaps marked in some way), because other projects on the chip may be OK. The quality of wafer fabrication for MPC79 is excellent, relative to the value of λ, and so this optical inspection may not be necessary for small projects.

When selecting diced chips for further packaging, note whether or not they are overglassed. Chips that are overglassed and have their package-cavities covered are well protected from possible contaminants, and can be placed into long term use. A small fraction of the wafers have been overglassed, so that any projects that function correctly can be placed into such long term use. However, even non-overglassed chips have relatively long-term functional lifetimes (a number of years), if they are kept boxed or covered and kept away from exposure to humid, salty atmospheres. The overglassing of chips sometimes visually obscures the details of the fabricated circuitry. Thus most of the MPC79 wafers will not be overglassed, so that projects can be more clearly studied and photographed through microscopes. In addition, non-overglassed chips can be internally probed if necessary for functional debugging. So, you're probably better off initially with non-overglassed chips. If your project works, you might have some overglassed chips packaged-up later.

Using the microscope to look at and show off your chip

It is always thrilling to take your first look through a microscope at a newly-implemented project that you have designed. It's almost magical to see, printed in fine detail on silicon, the project you spent all those long hours working on. If you want to see all the details of your project and how well the process worked, you will need a high-power microscope. This may help you to uncover various design rule errors, etc., that you hadn't noticed before (note that many chips will work even with slight technical design-rule errors, since the value of λ was conservative). Use unpackaged chips for this purpose, so as not to take a chance of smashing the bonding wires of your packaged chips. We strongly recommend that you also take a look at your chip using relatively low-powered stereo microscopes. Although you can't see fine detail with these, they provide more of an overview of project chips, and are a good way to show them off and explain them to others. You can also use these to confirm that the wire-bonding of your chip was properly done. You can obtain nice interference colors when using stereo microscopes, or even hand magnifiers, by shining light at a moderately low angle across the chip's surface, and rotating the chip very slowly until it is in a position where the colors appear. Under the right conditions, the circuitry will appear to "glow with color" and you will be able to see all the different layers, in different colors!

3. Project Locator Maps and List of Designers and Their Projects:

This section contains:

(i) a map for each of the two wafer types, giving the locations of the various die-types in the wafers.

(ii) a set of maps of the 12 project die-types, showing and encoding the positions of all 82 projects.

(iii) a list associating the projects' wafer/die/location-codes and the MPC79 file-system project-ID's.

(iv) a list of all 82 projects, identifying their designers and the functions of the projects.

The list of all 82 MPC79 projects groups the projects by university, listing them alphabetically by project ID within each university grouping. Each project entry in the list includes (i) the wafer/die/position-number code for the project, (ii) the project ID, (iii) the project designer(s), (iv) a short description of the project's function, (v) the project's bounding box dimensions in microns, and (vi) the area of the project in square mm.

Using the information in these maps and lists, participants can locate their project by determining the die-type (B, C, D, E, F, G, I, J, K, L, M, N) and then the project-number suffix for their project. The boxed, packaged chips returned to the university project-lab coordinators are marked with these die-type/project-number codes. The associated marked-up custom wire-bonding maps for each project are marked with the same code. When searching through diced but unpackaged chips, note that each chip contains a large die-type code-letter in its upper right-hand corner (see the Bonding Map section; the maps clearly show the code-letters).

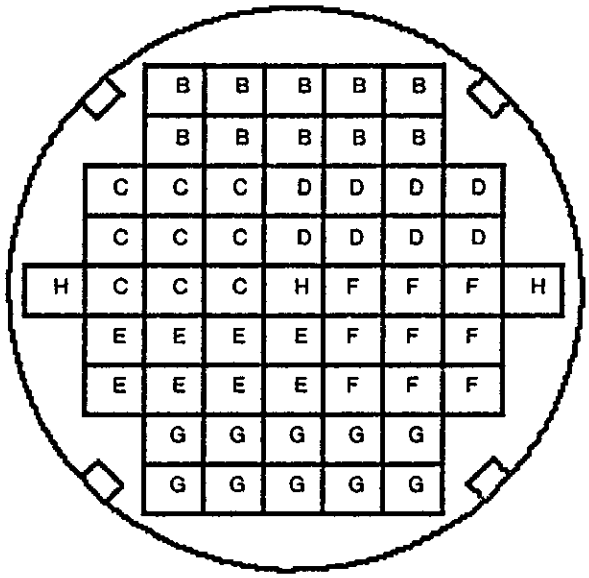

MPC79A
Wafer Map

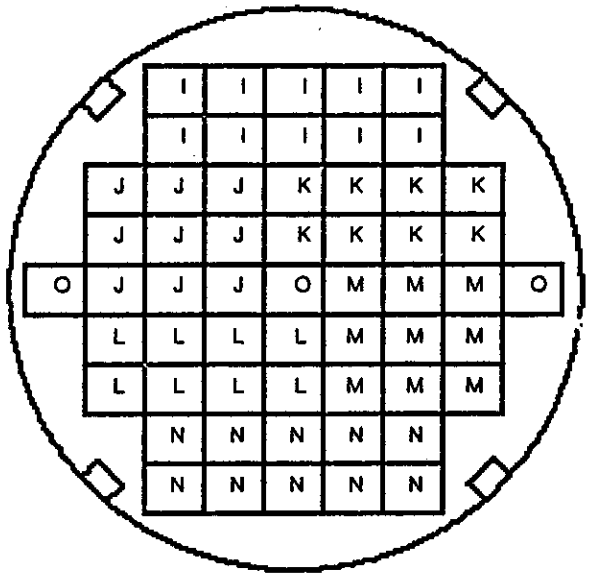

MPC79B
Wafer Map

AB

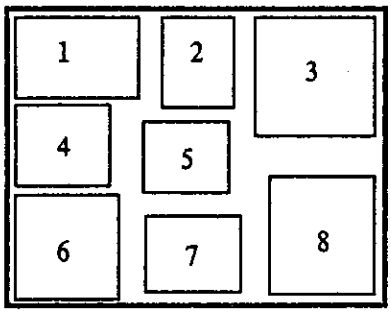

AC

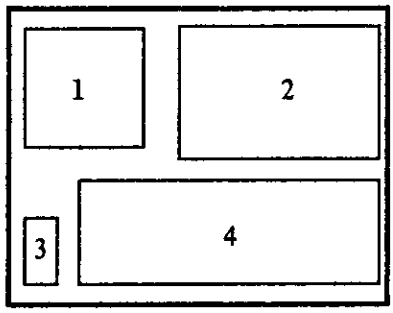

AD

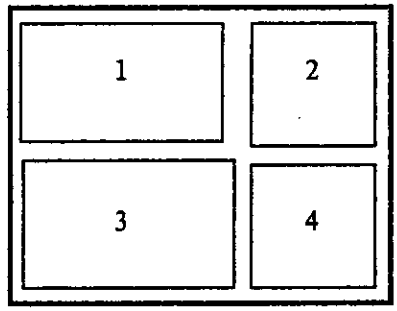

AE

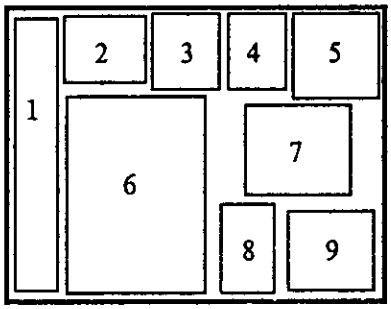

AF

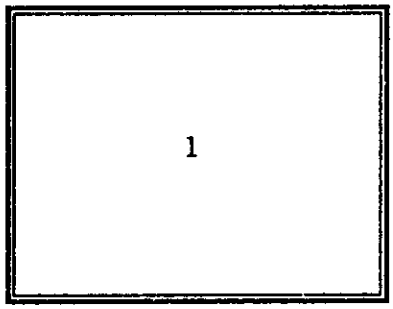

AG

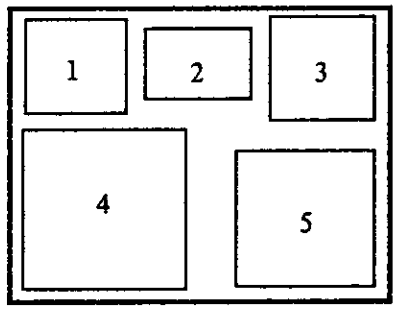

BI

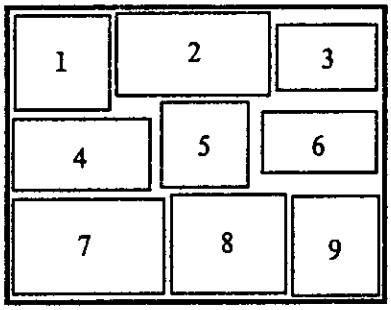

BJ

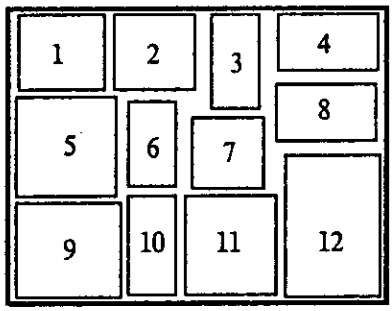

BK

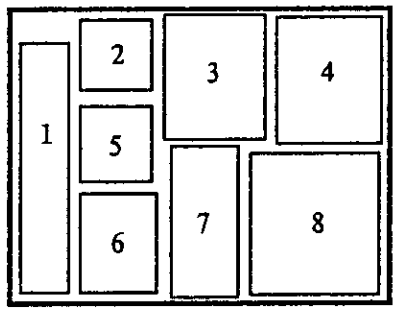

BL

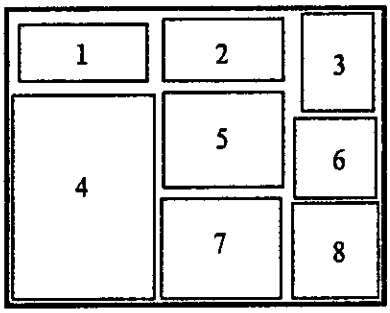

BM

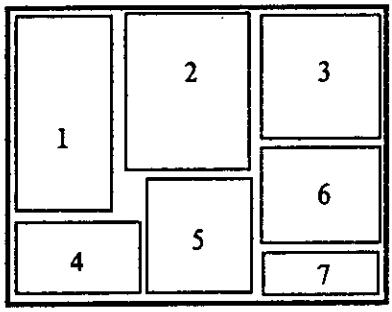

BN

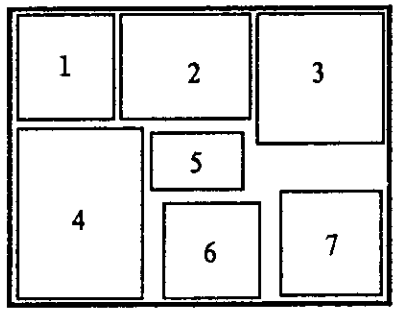

MPC79 Multiproject Chip Layouts

List of wafer-die-project codes and corresponding project ID's

Wafer MPC79A

AB-1 BataliMIT
AB-2 GramlichMIT
AB-3 FichtenbaumMIT
AB-4 KhouryMIT
AB-5 GoodrichMIT
AB-6 GrondalskiMIT
AB-7 PicardMIT
AB-8 AllenMIT

AC-1 HamiltonMIT
AC-2 PasemanMIT
AC-3 GlasserOT
AC-4 ChuMIT

AD-1 LuhukayUI
AD-2 HanesUI
AD-3 AdrianUI
AD-4 MontoyeUI

AE-1 GuptaCMU
AE-2 ClassUI
AE-3 MurrayOT
AE-4 RogersOT
AE-5 EbelingCMU
AE-6 KungCMU
AE-7 SongCMU
AE-8 HoeyCMU
AE-9 KehlOT

AF-1 Schip2

AG-1 WalpCT
AG-2 KathailMIT
AG-3 RivestMIT
AG-4 SnyderOT
AG-5 GoddeauMIT

AH-1 LHDocl
AH-2 RHDocl

Wafer MPC79B

BI-1 MacomberSU
BI-2 GehlbachSU
BI-3 MarkeeSU
BI-4 NoiceSU
BI-5 ElahianSU
BI-6 AtlasSU
BI-7 HerndonSU
BI-8 HannahSU
BI-9 WulfISU

BJ-1 CampbellCT
BJ-2 FuCT
BJ-3 PapachCT
BJ-4 LiCT
BJ-5 BartonCT
BJ-6 CocconiCT
BJ-7 PursifullCT
BJ-8 BozzutoCT
BJ-9 KingsleyCT
BJ-10 HoCT
BJ-11 WhitneyCT
BJ-12 TannerCT

BK-1 MathewsSU
BK-2 ZarghanSU
BK-3 FrolikSU
BK-4 BaskettSU
BK-5 Clark2SU
BK-6 OhChinSU
BK-7 BechtolsheimSU
BK-8 ClarkSU

BL-1 HellerCT
BL-2 EatonCT
BL-3 WatteyneCT
BL-4 MostellerCT
BL-5 GrayCT
BL-6 PinesCT
BL-7 DerbyCT
BL-8 PedersenCT

BM-1 LigockiCT
BM-2 DecuirUCB
BM-3 FungUCB
BM-4 LandmanUCB
BM-5 RumphCT
BM-6 EllisCT
BM-7 SequinUCB

BN-1 WatanabeUR
BN-2 LyonsUR
BN-3 KedemUR
BN-4 SohmUR
BN-5 TiloveUR
BN-6 UttSU
BN-7 TarsiSU

BO-1 LHDocl
BO-2 RHDocl

List of Designers and their Projects

CALTECH:

[Summary of designs from CalTech, updated 4-Dec-79 23:13:17]

BJ-5 BartonCT
Designer: Eric Barton
Description: LED array driver
Reserved space = 2126 x 2126 microns, Area = 4.52 sq mm

BJ-8 BozzutoCT
Designer: Rick Bozzuto
Description: Pulse width to binary converter
Reserved space = 2120 x 1288 microns, Area = 2.73 sq mm

BJ-1 CampbellCT
Designer: James Campbell
Description: Logical processing unit with internal registers
Reserved space = 1856 x 1704 microns, Area = 3.16 sq mm

BJ-6 CocconiCT
Designer: Alan Cocconi
Description: Array processor
Reserved space = 1896 x 1074 microns, Area = 2.04 sq mm

BL-7 DerbyCT
Designer: Howard Derby
Description: Associative Memory
Reserved space = 2170 x 2566 microns, Area = 5.57 sq mm

BL-2 EatonCT
Designer: Steve Eaton
Description: Counter/adder
Reserved space = 2500 x 1376 microns, Area = 3.44 sq mm

BM-6 EllisCT
Designer: Mike Ellis
Description: Stepping motor controller
Reserved space = 2000 x 2500 microns, Area = 5.00 sq mm

BJ-2 FuCT
Designer: Sai Wai Fu
Description: Square root generator
Reserved space = 1750 x 1626 microns, Area = 2.85 sq mm

BL-5 GrayCT
Designer: Moshe Gray
Description: Array processor
Reserved space = 2534 x 2082 microns, Area = 5.28 sq mm

BL-1 HellerCT
Designer: Jack Heller
Description: Digital filter
Reserved space = 2708 x 1326 microns, Area = 3.59 sq mm

BJ-10 HoCT
Designer: Kuo Ting Ho
Description: 10 bit rate multipler
Reserved space = 2120 x 1110 microns, Area = 2.35 sq mm

BJ-9 KingsleyCT
Designer: Chris Kingsley
Description: Serial Multipler
Reserved space = 2200 x 2064 microns, Area = 4.54 sq mm

CALTECH (cont.):

BJ-4 LiCT
Designer: Peggy Pey-Yun Li
Description: Two's-complement pipeline multiplier
Reserved space = 2176 x 1326 microns, Area = 2.89 sq mm

BM-1 LigockiCT
Designer: Terry Ligocki
Description: Scan converter chip
Reserved space = 2000 x 4108 microns, Area = 8.22 sq mm

BL-4 MostellerCT
Designers: Rick Mosteller, Greg Eflan, DIck Lang
Description: Stack-oriented micrprocessor
Reserved space = 4300 x 2996 microns, Area = 12.88 sq mm

BJ-3 PapachCT
Designer: A.C. Papachristidis
Description: Magnitude comparator
Reserved space = 2000 x 1126 microns, Area = 2.25 sq mm

BL-8 PedersenCT
Designer: Bruce Pedersen
Description: Asynchronous FIFO
Reserved space = 1896 x 2000 microns, Area = 3.79 sq mm

BL-6 PinesCT
Designer: Elliot Pines
Description: Expandable clocking pattern generator chip
Reserved space = 1780 x 1780 microns, Area = 3.17 sq mm

BJ-7 PursifullCT
Designer: Ralph Pursiful
Description: Self-Timed Queue
Reserved space = 1590 x 1590 microns, Area = 2.53 sq mm

BM-5 RumphCT
Designer: David Rumph
Description: DMA controller
Reserved space = 2442 x 2242 microns, Area = 5.47 sq mm

BJ-12 TannerCT
Designers: John Tanner and Richard Segal
Description: Single wire interface for a Manipulator (SWIM)
Reserved space = 2000 x 3000 microns, Area = 6.00 sq mm

AG-1 WalpCT
Designer: Pat Walp
Description: Array processor
Reserved space = 2126 x 2050 microns, Area = 4.36 sq mm

BL-3 WatteyneCT
Designers: Thierry Watteyne and Martine Savalle
Description: BCD/binary comparator
Reserved space = 2100 x 1600 microns, Area = 3.36 sq mm

BJ-11 WhitneyCT
Designer: Telle Whitney
Description: Address translator
Reserved space = 1940 x 2126 microns, Area = 4.12 sq mm

Carnegie-Mellon University:

[Summary of designs from CMU, updated 4-Dec-79 23:13:17]

AE-5 EbelingCMU
Designer: Carl Ebeling
Description: Rebound Sorter
Reserved space = 1856 x 1856 microns, Area = 3.44 sq mm

AE-1 GuptaCMU
Designer: Satish Gupta
Description: Video Buffer
Reserved space = 1006 x 5668 microns, Area = 5.70 sq mm

AE-8 HoeyCMU
Designer: Dan Hoey
Description: Experimental Adder
Reserved space = 1188 x 1976 microns, Area = 2.35 sq mm

AE-6 KungCMU
Designers: H. T. Kung, S. W. Song
Description: Image Processing Chip
Reserved space = 4160 x 2948 microns, Area = 12.26 sq mm

AE-7 SongCMU
Designer: Siang W Song
Description: A small database machine
Reserved space = 2224 x 1954 microns, Area = 4.35 sq mm

MIT:

[Summary of designs from MIT, updated 4-Dec-79 23:13:17]

AB-8 AllenMIT
Designers: Don Allen, Jerry Burchfiel
Description: Variable Length Field Decoder
Reserved space = 2218 x 2484 microns, Area = 5.51 sq mm

AB-1 BataliMIT
Designer: John Batali
Description: Zero-Crossing Detector for Image Processing
Reserved space = 2644 x 1738 microns, Area = 4.60 sq mm

AC-4 ChuMIT
Designers: Tam-Anh Chu, Nhi-Anh Chu, Steve McCormick
Description: Second order digital filter stage
Reserved space = 6146 x 2278 microns, Area = 14.00 sq mm

AB-3 FichtenbaumMIT
Designer: Matt Fichtenbaum
Description: A digital pulse rate monitor
Reserved space = 2500 x 2500 microns, Area = 6.25 sq mm

AG-5 GoddeauMIT
Designers: David Goddeau, Jonathan Sieber, Chris Terman
Description: A first-in, priority-out buffer
Reserved space = 2928 x 2954 microns, Area = 8.65 sq mm

MIT (cont.):

AB-5 GoodrichMIT
Designer: Earl Goodrich
Description: CRT controller
Reserved space = 1856 x 1520 microns, Area = 2.82 sq mm

AB-2 GramlichMIT
Designers: Wayne Gramlich, Carl Seaquist
Description: A writable PLA in which the programming of the AND and OR planes is defined by contents of static RAM cells. Also can program feedback loops to form finite state machines.
Reserved space = 1524 x 1906 microns, Area = 2.90 sq mm

AB-6 GrondalskiMIT
Designer: Robert Grondalski
Description: Writeable PLA
Reserved space = 2200 x 2200 microns, Area = 4.84 sq mm

AC-1 HamiltonMIT
Designer: Brian Hamilton
Description: Digital Alarm Clock
Reserved space = 2500 x 2500 microns, Area = 6.25 sq mm

AG-2 KathailMIT
Designers: Vinod Kathail, Keshav Pingali
Description: an interpreter for mapping programs onto a data flow computer
Reserved space = 1590 x 2228 microns, Area = 3.54 sq mm

AB-4 KhouryMIT
Designer: John Khoury
Description: Up-Down counter with programmable modulus
Reserved space = 2000 x 1726 microns, Area = 3.45 sq mm

AC-2 PasemanMIT
Designer: Bill Paseman
Description: Music Synthesizer
Reserved space = 4126 x 2842 microns, Area = 11.73 sq mm

AB-7 PicardMIT
Designer: Len Picard
Description: Variable format field extractor and compactor
Reserved space = 2000 x 1688 microns, Area = 3.38 sq mm

AG-3 RivestMIT
Designers: Ron Rivest, Len Adleman, Adi Shamir
Description: Section of a Multiplier
Reserved space = 2250 x 2250 microns, Area = 5.06 sq mm

Stanford University:

[Summary of designs from Stanford University, updated 4-Dec-79 23:13:17]

BI-6 AtlasSU
Designers: Les Atlas, Doug Galbraith
Description: This project is an neural-stim. interval timer
Reserved space = 2478 x 1378 microns, Area = 3.41 sq mm

BK-4 BaskettSU
Designer: Forest Baskett
Description: This project is an Ethernet synchronizer
Reserved space = 2240 x 2720 microns, Area = 6.09 sq mm

BK-7 BechtolsheimSU
Designers: Andy Bechtolsheim, Thomas Gross
Description: A parallel search table for log arithmetic
Reserved space = 1514 x 3180 microns, Area = 4.81 sq mm

BK-5 Clark2SU
Designer: Jim Clark
Description: This project is a self-timed clock element
Reserved space = 1606 x 1688 microns, Area = 2.71 sq mm

BK-8 ClarkSU
Designer: Jim Clark
Description: This project is a simple graphics ALU
Reserved space = 2976 x 2764 microns, Area = 8.23 sq mm

BI-5 ElahianSU
Designers: Kamran Elahian, Fred Basham
Description: This project is a UART line speed determiner
Reserved space = 1856 x 1856 microns, Area = 3.44 sq mm

BK-3 FrolikSU
Designers: Bill Frolik, Roderick Young
Description: This project is a digital timer
Reserved space = 2120 x 2684 microns, Area = 5.69 sq mm

BI-2 GehlbachSU
Designers: Steve Gehlbach, Joe Sharp, Bill Jansen
Description: This project is a fast 16-input adder
Reserved space = 3180 x 1856 microns, Area = 5.90 sq mm

BI-8 HannahSU
Designers: Peter Eichenberger, Marc Hannah
Description: This project is a rectangle generator
Reserved space = 2386 x 2140 microns, Area = 5.11 sq mm

BI-7 HerndonSU
Designers: Matt Herndon, Jeff Thorson
Description: This project is a typesetting machine
Reserved space = 3170 x 2000 microns, Area = 6.34 sq mm

BI-1 MacomberSU
Designers: Scott Macomber, Bob Clark
Description: This project is a parallel/serial multiplier
Reserved space = 2000 x 2000 microns, Area = 4.00 sq mm

BI-3 MarkeeSU
Designers: Pat Markee, Irene Watson
Description: This project is a digital clock
Reserved space = 2120 x 1424 microns, Area = 3.02 sq mm

BK-1 MathewsSU
Designers: Rob Mathews, John Newkirk
Desciption: This project is the infamous Buffalo chip
Reserved space = 5180 x 1134 microns, Area = 5.87 sq mm

BI-4 NoiceSU
Designers: David Noice, Neil Midkiff
Description: This project is a multiplier/divider
Reserved space = 2888 x 1576 microns, Area = 4.55 sq mm

BK-6 OhChinSU
Designers: Soo-Young Oh, Dae-Je Chin
Description: An automatic thermostat time controller
Reserved space = 2120 x 1700 microns, Area = 3.60 sq mm

BN-7 TarsiSU
Designers: Mike Tarsi, Nagatsugu Yamanouchi
Description: This project is a multifunction digital clock
Reserved space = 2140 x 2276 microns, Area = 4.87 sq mm

BN-6 UttSU
Designers: Steve Utt, Shalom Ackelsberg
Description: This project is part of a pancreas prosthesis
Reserved space = 2000 x 2000 microns, Area = 4.00 sq mm

BI-9 WulffSU
Designers: Bob Wulff, Tom Bennett
Description: This project is a bit slice of a multiplier
Reserved space = 2120 x 1856 microns, Area = 3.93 sq mm

BK-2 ZarghanSU
Designers: Bahman Zargham, Jerry Huck
Description: This project is a multiplexed communications link
Reserved space = 1590 x 1550 microns, Area = 2.46 sq mm

U.C.Berkeley:

[Summary of designs from U.C.Berkeley, updated 4-Dec-79 23:13:17]

BM-2 DecuirUCB
Designers: J. Decuir, C.H.Sequin
Description: Squareroot of 3 approximator for
radix-3 block in FFT computer
Reserved space = 2650 x 3278 microns, Area = 8.69 sq mm

BM-3 FungUCB
Designers: W.-C. Fung, C.H.Sequin
Description: General purpose barrel shifter for sraggered ,
pipelined data in an FFT computer
Reserved space = 2484 x 2650 microns, Area = 6.58 sq mm

BM-4 LandmanUCB
Designer: Howard A. Landman
Description: This project is a reprogrammable PLA, with
8 each inputs, pterms, and (tri-state) outputs.
Reserved space = 2600 x 1590 microns, Area = 4.13 sq mm

BM-7 SequinUCB
Designer: Carlo H. Sequin
Description: Dual 16-stage FIFO with double rail signalling
Reserved space = 2460 x 980 microns, Area = 2.41 sq mm

Univ. of Illinois:

[Summary of designs from University of Illinois, updated 4-Dec-79 23:13:17]

AD-3 AdrianUI
Designers: Frank Adrian, Nick Fiduccia, Bud Pflug
Description: Functional equivalent of AMD 2901 ALU
to compare MOS, TTL
Reserved space = 2710 x 4388 microns, Area = 11.89 sq mm

AE-2 ClassUI
Designers: Class
Description: Twos complement 4 x 4 array multiplier
Reserved space = 1714 x 1498 microns, Area = 2.57 sq mm

AD-2 HanesUI
Designers: Larry Hanes, Dave Yen
Description: Twos complement array divider
Reserved space = 2616 x 2636 microns, Area = 6.90 sq mm

AD-1 LuhukayUI
Designer: Joe Luhukay
Description: Pipelined multiplier, registers also used for testability
Reserved space = 2572 x 4140 microns, Area = 10.65 sq mm

AD-4 MontoyeUI
Designers: Bob Montoye, Al Casavant
Description: Carry lookahead adder
(soln. proposed by Gajski and Kung)
Reserved space = 2628 x 2626 microns, Area = 6.90 sq mm

Univ. of Rochester:

[Summary of designs from University of Rochester, updated 4-Dec-79 23:13:17]

BN-3 KedemUR
Designers: Gershon Kedem and Michel Denber
Description: Infinite precision multiplier
Reserved space = 2698 x 2786 microns, Area = 7.52 sq mm

BN-2 LyonsUR
Designer: Bob Lyons
Description: Programmable Frequency Generator
Reserved space = 2748 x 2276 microns, Area = 6.25 sq mm

BN-4 SohmUR
Designers: Larry Sohm, Pat Chan, Bill Notowitz
Description: Digital Phase lock loop
Reserved space = 3610 x 2634 microns, Area = 9.51 sq mm

BN-5 TiloveUR
Designers: Bob Tilove, Jarek Rossignac
Description: This is a bit slice coordinate transformer
Reserved space = 1934 x 1326 microns, Area = 2.56 sq mm

BN-1 WatanabeUR
Designer: Yuki Watanabe
Description: Sorter slice
Reserved space = 2008 x 2240 microns, Area = 4.50 sq mm

Other places:

[Summary of designs from Other places, updated 4-Dec-79 23:13:17]

AC-3 GlasserOT
Designer: Lance Glasser, MIT, via Univ. of Washington
Description: Modulo-6 counter for dice game
Reserved space = 1486 x 808 microns, Area = 1.20 sq mm

AE-9 KehlOT
Designers: Ted Kehl, Ram Rao, Ed Lazowska,
Univ. of Washington, Seattle
Description: Address intercept logic for microcomputer
Reserved space = 1818 x 1782 microns, Area = 3.24 sq mm

AE-3 MurrayOT
Designer: John Murray, Univ. of Colorado, Colorado Springs,
via Univ. of Washington
Description: 3-bit identity comparator
Reserved space = 1512 x 1642 microns, Area = 2.48 sq mm

AE-4 RogersOT
Designer: Mike Rogers, Univ. of Bristol, Bristol, England,
via Univ. of Washington
Description: Simple 3-bit enciphering/deciphering chip.
Reserved space = 1248 x 1708 microns, Area = 2.13 sq mm

AF-1 Schip2
Designers: Gerry Sussman, Jack Holloway, Guy Steele, Alan Bell
MIT-AI Laboratory/Xerox PARC-SSL
Description: Lisp Microprocesser
Reserved space = 5926 x 7548 microns, Area = 44.73 sq mm

AG-4 SnyderOT
Designer: Larry Snyder, Yale University,
via University of Washington
Description: A binary tree processor that computes boolean
functions, with inputs at the leaves and output at the root.
Reserved space = 3418 x 3430 microns, Area = 11.72 sq mm

AH-1 LhDoc1
Designer: Lynn Conway
Description: This is the Left Half of a "document chip",
describing MPC79, for use on MPC79A wafers.
Reserved space = 2918 x 4688 microns, Area = 13.68 sq mm

BO-1 LhDoc2
This is the Left Half of the "document chip",
for use on MPC79B wafers.
Reserved space = 2918 x 4688 microns, Area = 13.68 sq mm

AH-2 RHDoc1
Designer: Lynn Conway
Description: This is the Right Half of a "document chip",
flowcharting MPC79, for use on MPC79A wafers.
Reserved space = 3548 x 4424 microns, Area = 15.70 sq mm

BO-2 RHDoc2
This is the Right Half of the "document chip",
for use on MPC79B wafers.
Reserved space = 3548 x 4424 microns, Area = 15.70 sq mm

4. Starting Frame Documentation

The following pages describe the structure, function, and use of the various objects in the MPC79 starting frame (a starting frame is the set of all those layout artifacts, not associated with any particular design project, which serve to convey a project chip through maskmaking, wafer fabrication, and electrical testing). Included are the die dimensions and layout, the scribe-line profile, and descriptions of the alignment marks, layer codes, critical dimension testers, etch test patterns, die identification codes, discrete test transistors, ring oscillator, test resistors, contact tester, layer tester, and die code-letters.

Inter-Office Memorandum

To	MPC79 Distribution	Date	December 17, 1979
From	Rich Pasco	Location	Palo Alto
Subject	MPC79 Starting Frame	Organization	PARC-SSL-LSI

XEROX

Filed on: [ivy]<Pasco>MPC79>StartingFrameMemo.Press

This memo describes the Starting Frame used in generation of the MPC79 chip set. The Starting Frame is a collection of items common to all chips in the set. It consists of a scribe line to make it easier to cut the wafer into dice, and a top strip containing alignment marks, lithography test patterns, and test devices.

Die Dimensions and Layout

Overall Die Dimensions: 7696 x 6477 microns

The dimensions are:

Vertical:		Horizontal:	
Scribe-Line Profile:	66 microns.	Scribe-Line Profile:	66 microns.
Setback:	21 microns.	Setback:	8 microns.
Top-Strip:	370 microns.	Maximum project:	7548 microns.
Setback:	20 microns.	Setback:	8 microns.
Maximum project:	5926 microns.	Scribe-Line Profile:	66 microns.
Setback:	8 microns.		
Scribe-Line Profile:	66 microns.		
Total:	6477 Microns (=255 Mils)	Total:	7696 Microns (~303 Mils)

Scribe-Line Profile.

The scribe-line profile includes the scribe line itself and a metal band. The scribe line itself is a coincident contact cut, diffusion, and overglass cut, 41 microns to either side of center. This removes all oxide, allowing the wafer saw to cut directly into the silicon wafer. 10 microns from the scribe line is a 15-micron-wide band of metal which serves as a visual guide to the saw operator and separates the project area from the saw area.

Alignment Marks.

Two kinds of alignment marks are provided. In the upper left corner of the die is a square appearing on all layers; this provides a gross alignment indicator. For fine alignment, to its right are a sequence of "Squares" and "Fortresses" as described in Section 6.1 of the Xerox PARC/SSL Report, *A Guide to LSI Implementation*, 2nd Edition, by Hon & Sequin (Sect. 7.1 in 1st Edition).

Layer Codes, Critical Dimension Testers, and Etch Test Patterns (Ells)

Below the alignment marks appears a set of six similar patterns, one for each of the six layers. Each consists of three items: the layer code, from {DIF, IMP, POL, CUT, MET, PAD}, a cross-shaped critical dimension tester, and a set of Ell-shaped etch test patterns.

Critical Dimension Testers are crosses with line widths as follows:

DIF	5.0 micron	CUT	5.0 micron
IMP	10.0 micron	MET	10.0 micron
POL	5.0 micron	PAD	10.0 micron

Etch test patterns (ells) are 65 microns high. Smaller ells, with 2.5 micron linewidth and spacing, are nested inside the upper right corner of larger ells with 5 micron linewidth and spacing. In addition, there is a vertical 5-micron-wide bar down the left side of the pattern to make measurement easier.

The polysilicon critical dimension tester and etch test pattern are repeated over thin oxide (diffusion mask) to the right of the alignment marks.

Although digital stretching of diffusion layer is specified for all active devices on MPC79, no stretching is performed on critical dimension testers and etch test patterns.

Identification Code

Each die has a unique seven-character identifier, e.g. "MPC79AF". The first five characters "MPC79" indicate that this chip is a 1979 MultiUniversity Project Chip. The next two characters indicate the wafer and die, respectively. There are fourteen die types (B through O) in MPC79 on two wafers (A and B); thus code letters are in {AB, AC, AD, AE, AF, AG, AH, BI, BJ, BK, BL, BM, BN, BO}. For example, letters "MPC79AB" indicate that this is die type B on wafer type A.

Discrete Test Transistors

Discrete devices are shown for measurement of dc parameters such as threshold voltages. Figure 1 shows the location within the starting frame and pad connections for the test devices.

Four devices, with sources and gates are bussed together, but having separate drains, are provided. All have 4 lambda (10 microns) wide by 2 lambda (5 microns) long channels. There is one device of each of the following types:

Metal gate, field oxide.
Poly gate, field oxide.
Poly gate, thin oxide.
Poly gate, thin oxide, Ion Implant (depletion mode),

Ring Oscillator

The ring oscillator is used to estimate the speed of devices made using this process. It consists of nineteen identical inverters in a circle, and a twentieth as a buffer to drive a line to a standard output pad, where there is more buffering to drive the outside world. Each inverter in the circle is minimum geometry, K=4. Pulldowns have 4 lambda (lambda = 2.5 microns) wide by 2 lambda long channels. Pullups have 2 lambda wide by 4 lambda long channels. Since the oscillator period T equals the delay twice around the loop, the inverter pair delay is T/n, where n is between 19 and 20 since eighteen of the inverters drive only one similar inverter but the nineteenth drives two.

Test Resistors

There are two Van Der Pauw structures, one in polysilicon and one in diffusion, with identical geometry as shown in Figure 2. In a Van Der Pauw structure, there are two pads at each end of each resistor to allow separate current feed and voltage sense paths to the layer being tested, so that probe and contact resistance does not introduce error into the measurement. These devices provide a long aspect ratio (266.5 to 280.5 squares, depending on how corners are treated*) path of minimum width (5 microns = 2 lambda). Because of edge effects, sheet resistivity estimates from these measurements (sheet resistivity = total resistance divided by aspect ratio) are only approximate for other than 2 lambda wide lines.

*For further information on paths that go around corners in sheets of resistive material, see D. Vitkovitch, Ed., *Field Analysis, Experimental and Computational Methods*, D. Van Nostrand Co., Ltd., 1966.

Contact Tester

The contact tester has 3 sections, each consisting of 100 Mead & Conway minimum geometry contacts in series: metal-diffusion, metal-poly, and diffusion-poly-butting, with pads at each end and between each section. There should be continuity between any pair of pads.

Layer Tester

This layer tester was designed by Rick Davies and described in Hon & Sequin, *A Guide to LSI Implementation*, Second Edition, Section 6.2.1 (was 7.2.1 in First Edition) A polysilicon snake zigzags across thin oxide (diffusion mask), forming a large transistor. (Channel 2095 lambda wide by 2 lambda long). On top of this, a metal snake zigzags between two interdigitated metal combs, crossing the poly snake many times.

There should be continuity between:

Metal Snake In and Metal Snake Out
Poly Snake In and Poly Snake Out
Diffusion Left End and Diffusion Right End (when Poly Snake is high)

There should be no continuity between:

Metal Snake and Top Metal Comb
Metal Snake and Bottom Metal Comb
Poly Snake and Top Metal Comb
Poly Snake and Bottom Metal Comb
Poly Snake and Metal Snake
Poly Snake and Diffusion (be careful - thin gate oxide)
Diffusion and any Metal feature
Diffusion Left End and Diffusion Right End (when Poly Snake is low)

Die Code Letter

The die code letter is repeated in the upper right corner of the die. Hopefully this will be legible to the unaided eye.

MPC79 Starting Frame Test Devices

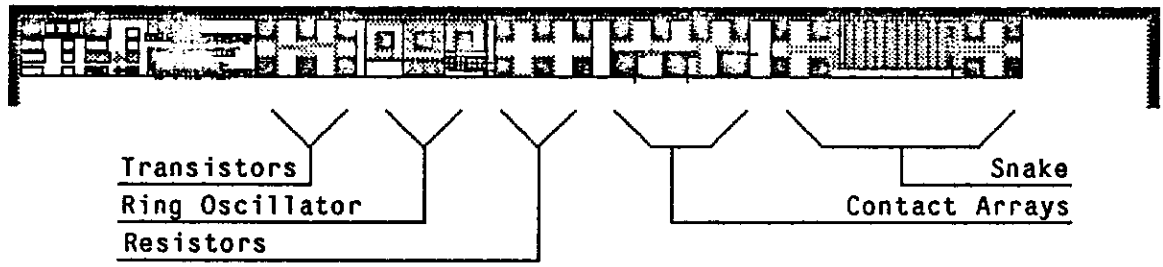

Transistors

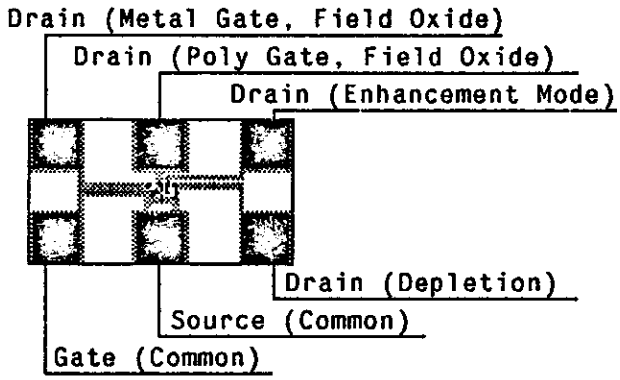

Ring Oscillator

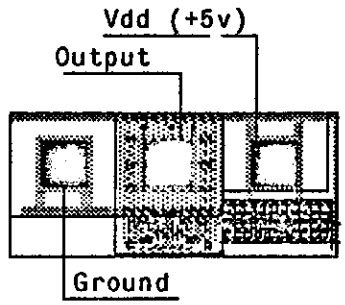

Resistors

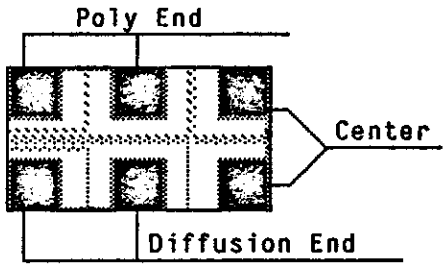

Contact Arrays

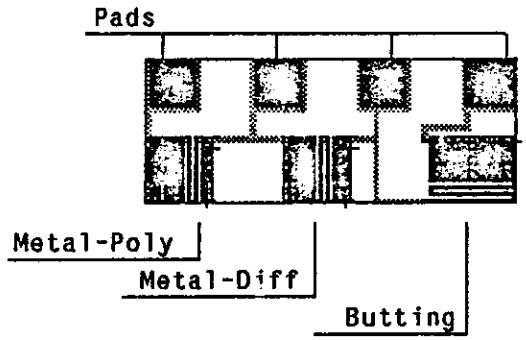

Snake

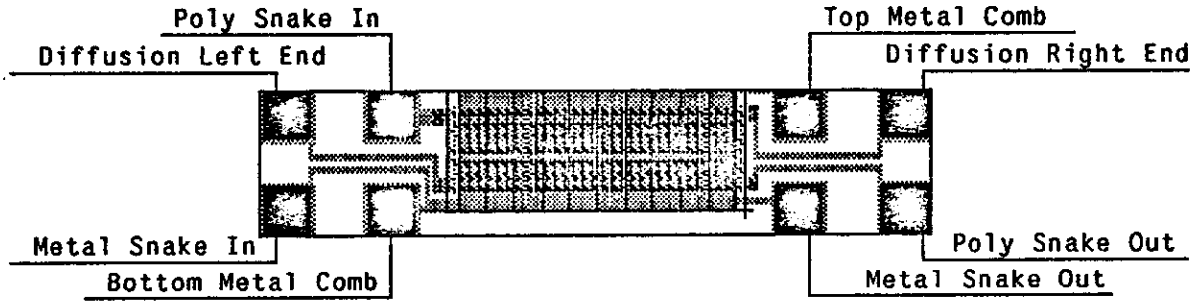

XEROX PARC	Project	Title	File	Designer	Rev	Date	Figure
	MPC79	Starting Frame Test Devices	TestDevices.Press	Pasco	A	12/11/79	1

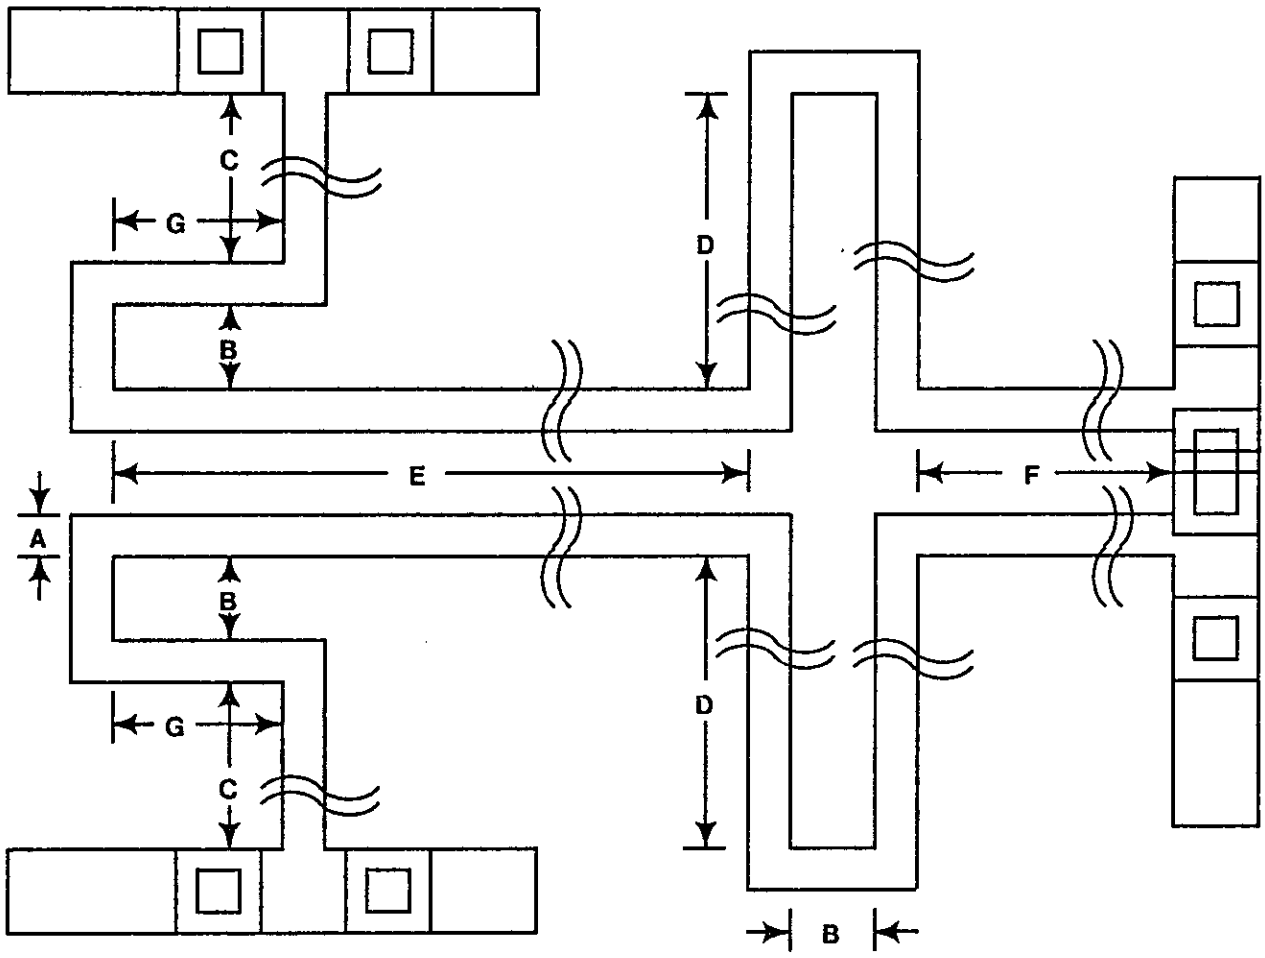

Dimension	Lambda	Microns
A	2	5
B	4	10
C	60	150
D	68	170
E	180	450
F	72	180
G	77	192.5
Inside Path Length*	533	1332.5
Outside Path Length*	561	1402.5

*Inside Path Length is shortest path, following inside edges.
Outside Path Length is path following outside edges.

XEROX PARC	Project	Title	File	Designer	Rev	Date	Figure
	MPC79	Test Resistors	Resistors.Sil	Pasco	A	12/11/79	2

5. Electrical & Process Measurements

This section provides the results of electrical measurements made at the Hewlett-Packard Integrated Circuit Processing Laboratory using the MPC79 starting frame test structures. Included are the ring oscillator frequency as a function of VDD, the various threshold voltages, certain electrical parameters, and basic information about the process, such as oxide thicknesses.

The transit time, τ, of minimum-sized transistors can be derived from this information, and then used by designers to estimate the maximum clock frequencies for their projects. The nineteen-stage ring oscillator "rings" at ~ 17 MHz (at VDD = 5v). Thus the inverter-pair-delay in the oscillator equals 3 ns (see Section 4, Starting Frame Documentation). Let's assume that the effective fanout, f, including parasitics is approximately equal to 2. The inverter-pair-delay $= f(k+1)\tau = 3$ ns. Therefore, we find that the transit time τ is approximately 0.3 ns.

HEWLETT hp PACKARD

3500 Deer Creek Road, Palo Alto, California 94304, U.S.A., Telephone 415 494-1444, TWX 910 373 1267

FROM: Mike Beaver DATE: December 31, 1979

TO: MPC 79 Distribution SUBJECT: MPC 79 Wafer Data

The attached data on the initial two runs for MPC79 represents in-process data from test wafers generated in the normal course of wafer fabrication plus electrical measurements taken on one of the wafers from each lot. This data is enclosed to give you some benchmark data against which you can compare your actual test results. Most of the electrical data was taken with a curve tracer and is subject to appropriate tolerances.

Additionally, wafer to wafer variations exist. If you need more precise data, you can generate it by probing the test devices on the border of each chip.

The follow-on run, from which you may receive additional chips will be similarly characterized. If you have specific questions about the processing, I will attempt to answer them.

Mike

WAFER DATA

	RUN KDEI1 Series A	RUN KDEI2 Series B
Thresholds		
Dep	- 4.9	- 4.1
Enh	1.0	0.8
Poly-gate Field	21.0	20.0
Metal-gate Field	21.0	20.0
Resistances		
Poly Resistor	6.4K	8.4K
Diff Resistor	3.6K	3.9K
Poly-Metal Contacts	2.3K	2.6K
Diff-Metal Contacts	1.4K	1.4K
Butting Contacts	3.6K	4.3K
Breakdown		
Diff-Substrate	34.0V	34.5V
Oscillator Frequency		
@ 5V	17MHz	17.7MHz
Process Parameters		
Gate Oxide Thickness	990Å	1000Å
Field Oxide Thickness	14600Å	14600Å
Poly Thickness	5370Å	5300Å
Intermediate Oxide Thickness	5000Å	5400Å
Diffusion Sheet Resistivity	17.2-17.8 Ω/□	16.2-17.1 Ω/□

6. Bonding Maps for the 12 Project Die Types

A marked-up wire-bonding map is included with every packaged chip being shipped to the project lab coordinators. These individually "customized", marked-up maps document the package pin-outs of the projects, so that project designers can determine how hook up the packages for functional testing of their projects.

In addition to the packaged chips, each of the project lab coordinators is also being sent a selection of diced but unmounted chips. The following section contains a full set of *unmarked* wire-bonding maps for the 12 different project die-types. Copies of these maps may be marked-up to document the pin-outs of any of the unmounted chips that are packaged later on (these will prove especially useful if later packaging leads to different pin-outs than for the original group of chips).

These maps also convey some notion of the scope and complexity of the various MPC79 designs, and of the overall MPC79 effort.

MPC79 A8

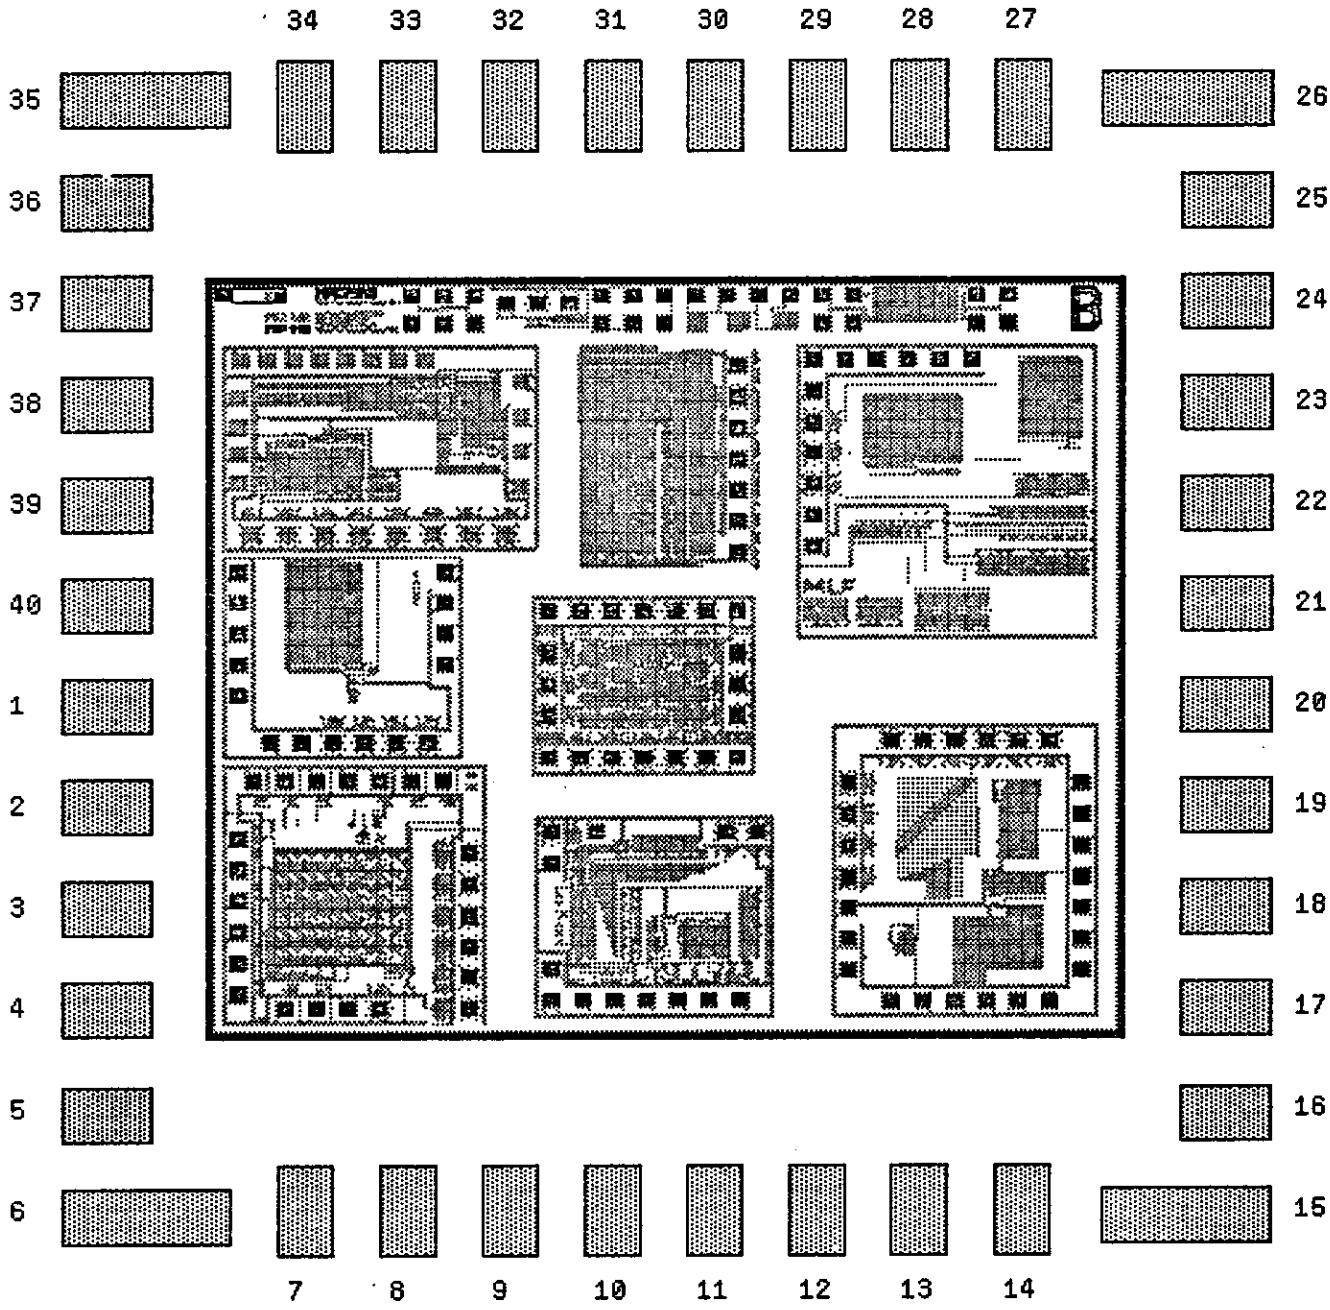
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 AC

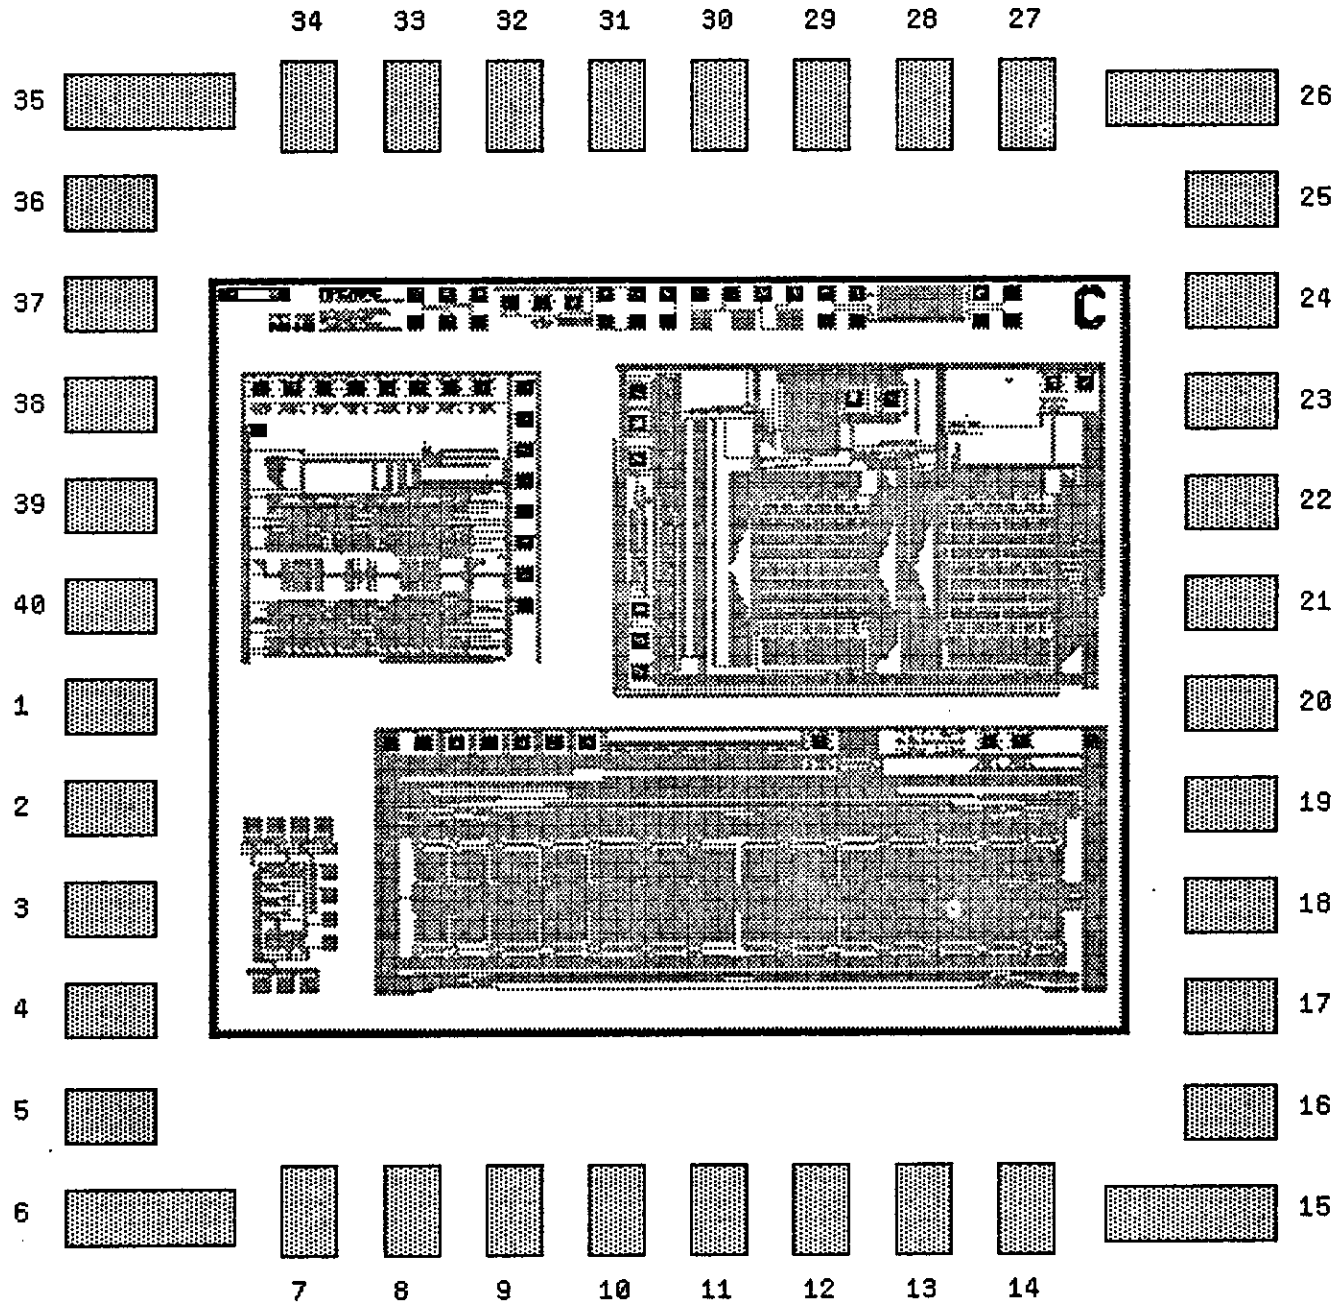
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 AD

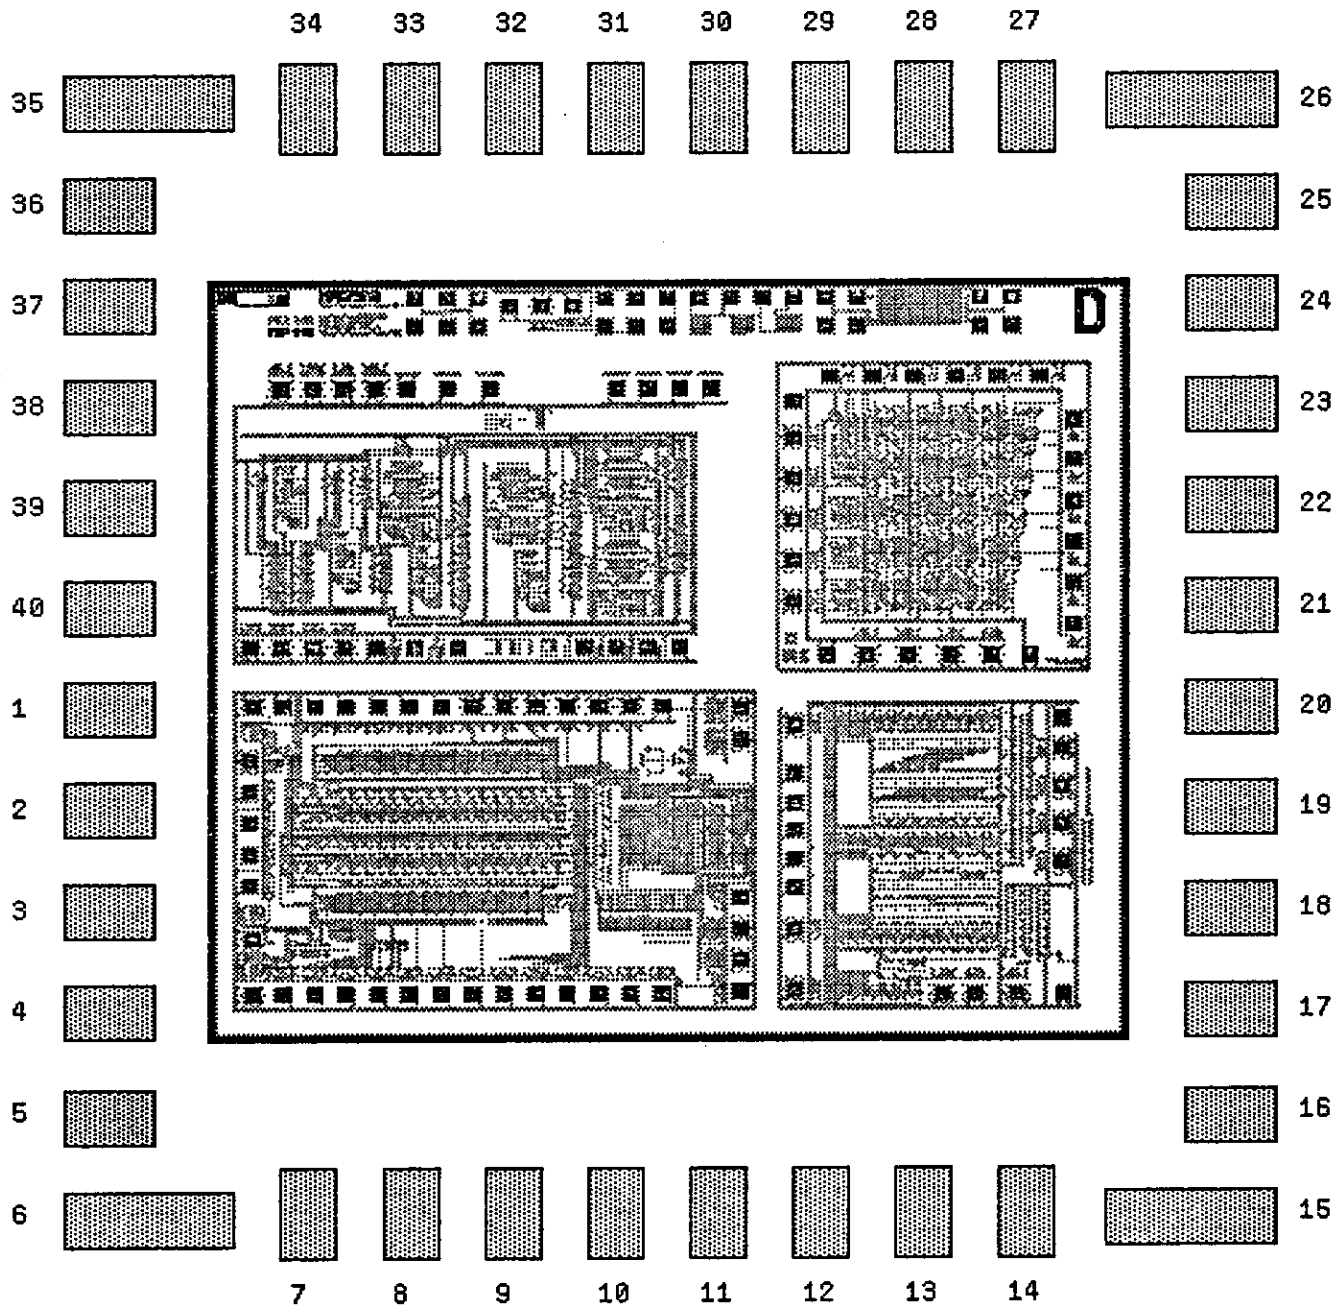
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14
D

MPC79 AE

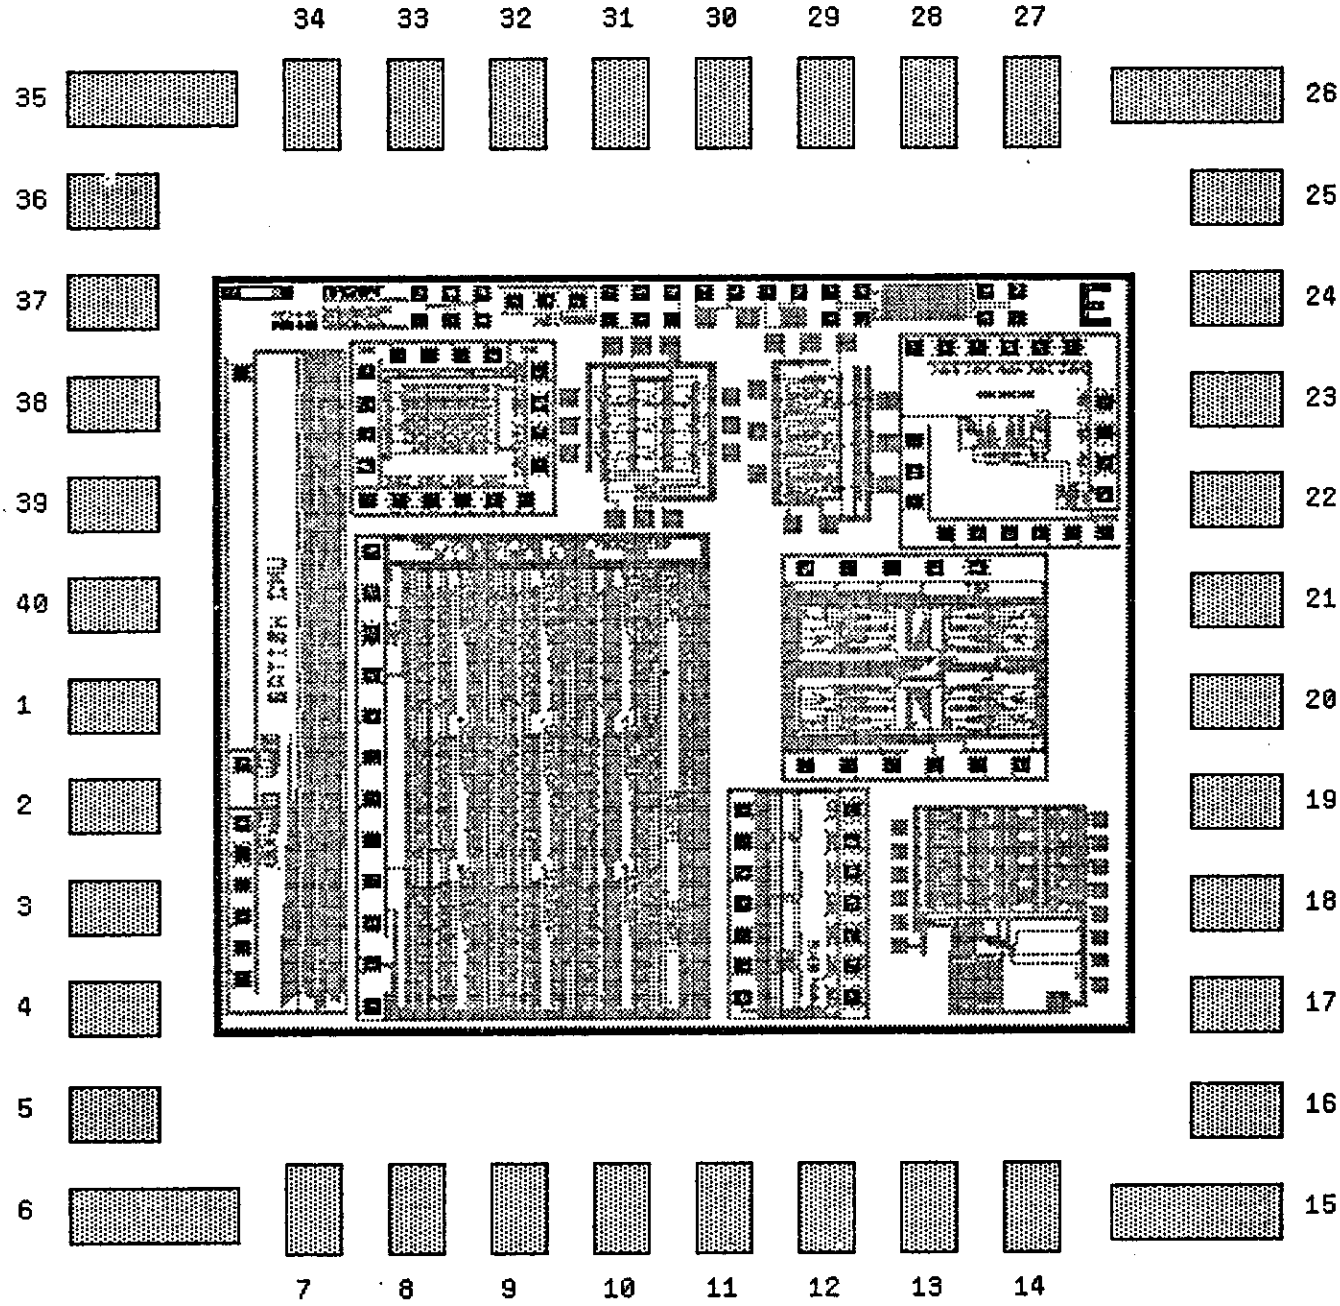
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 AF

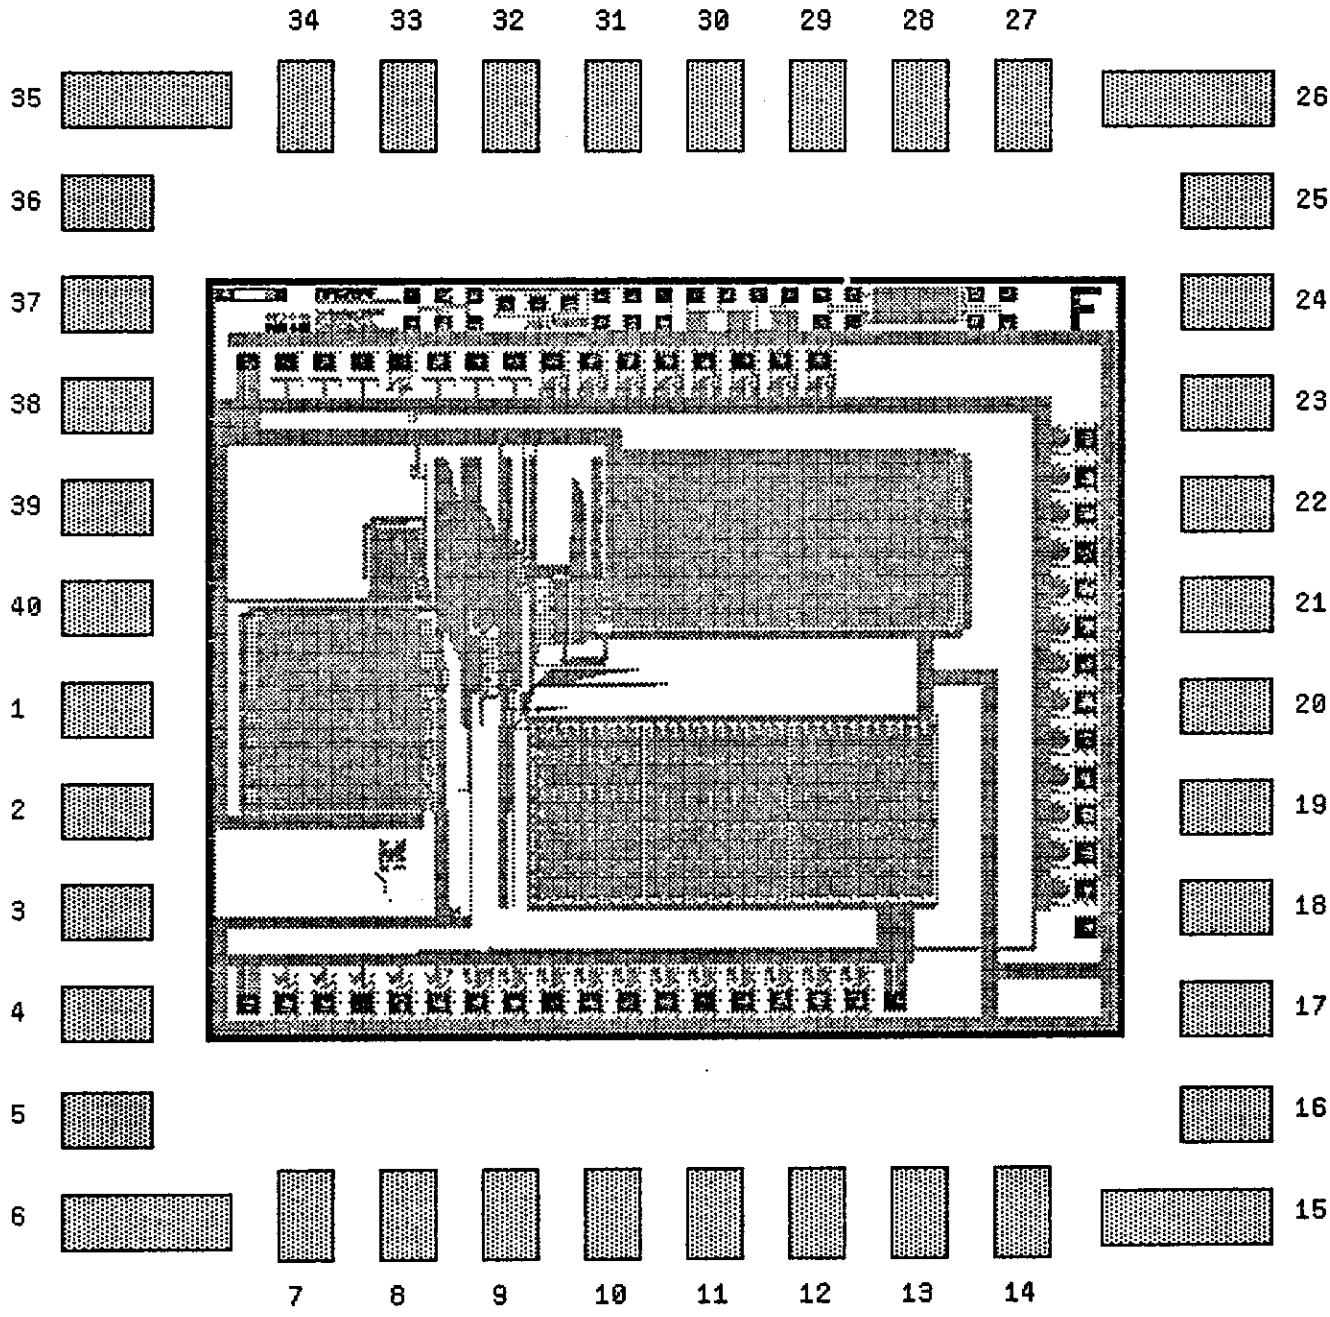
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 AG

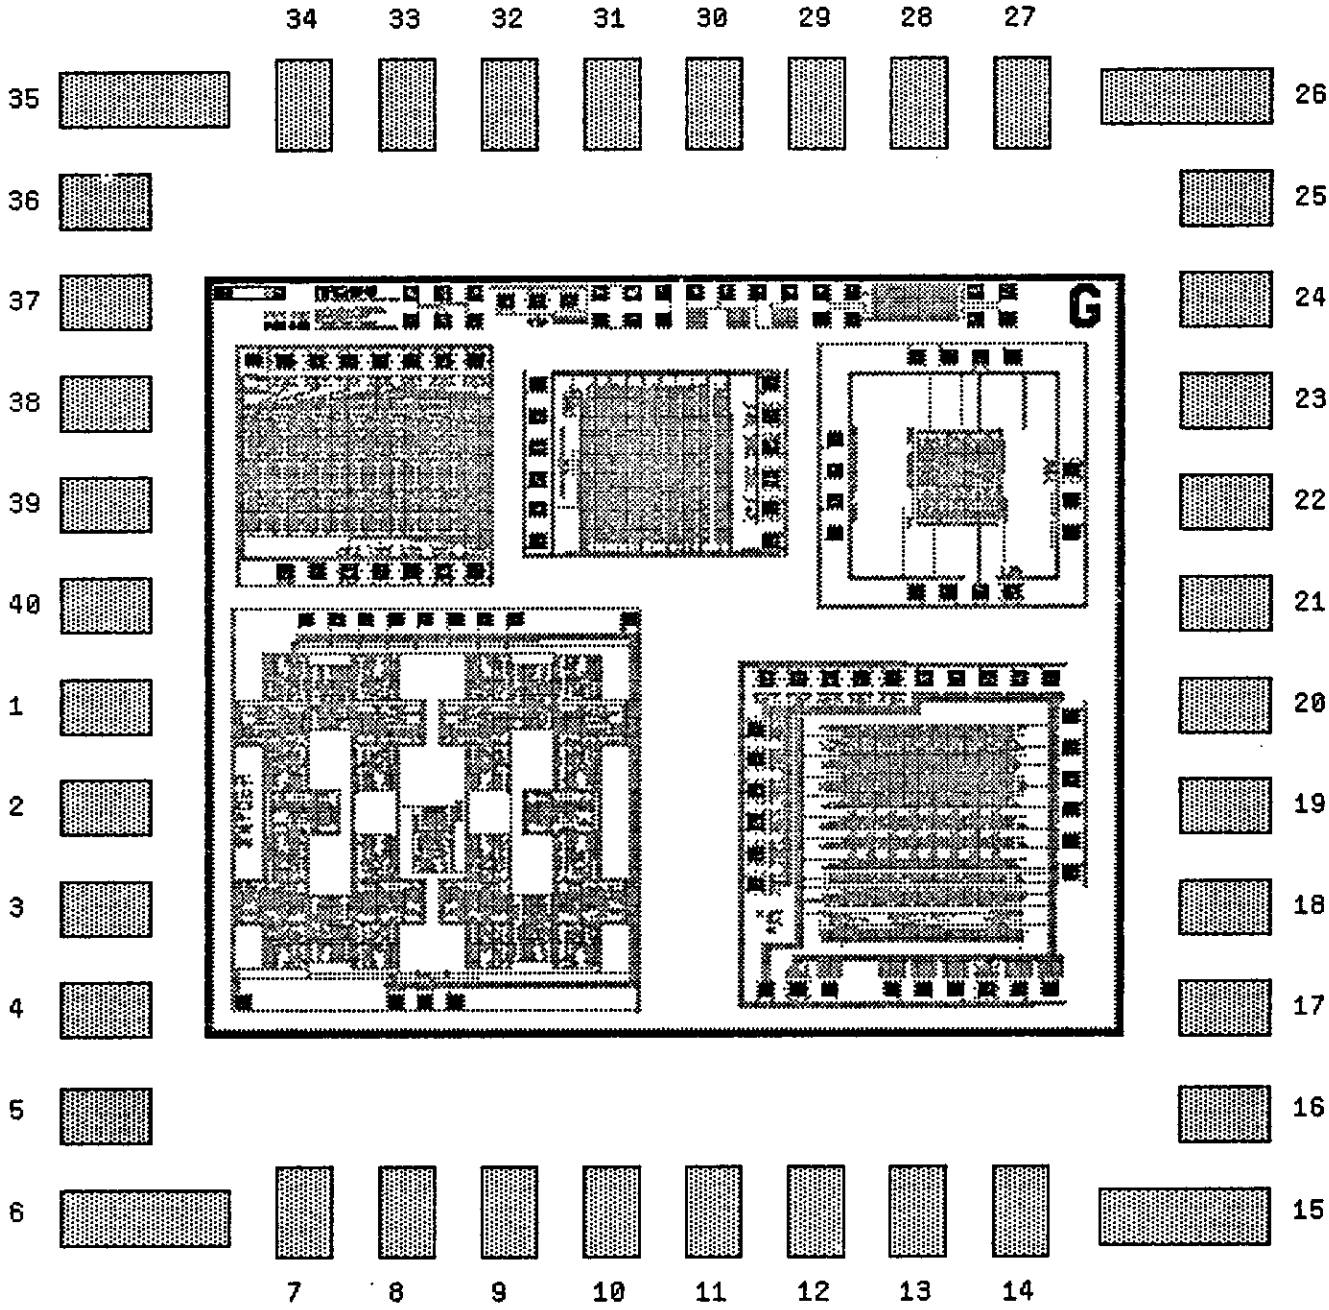
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14
G

MPC79 BI

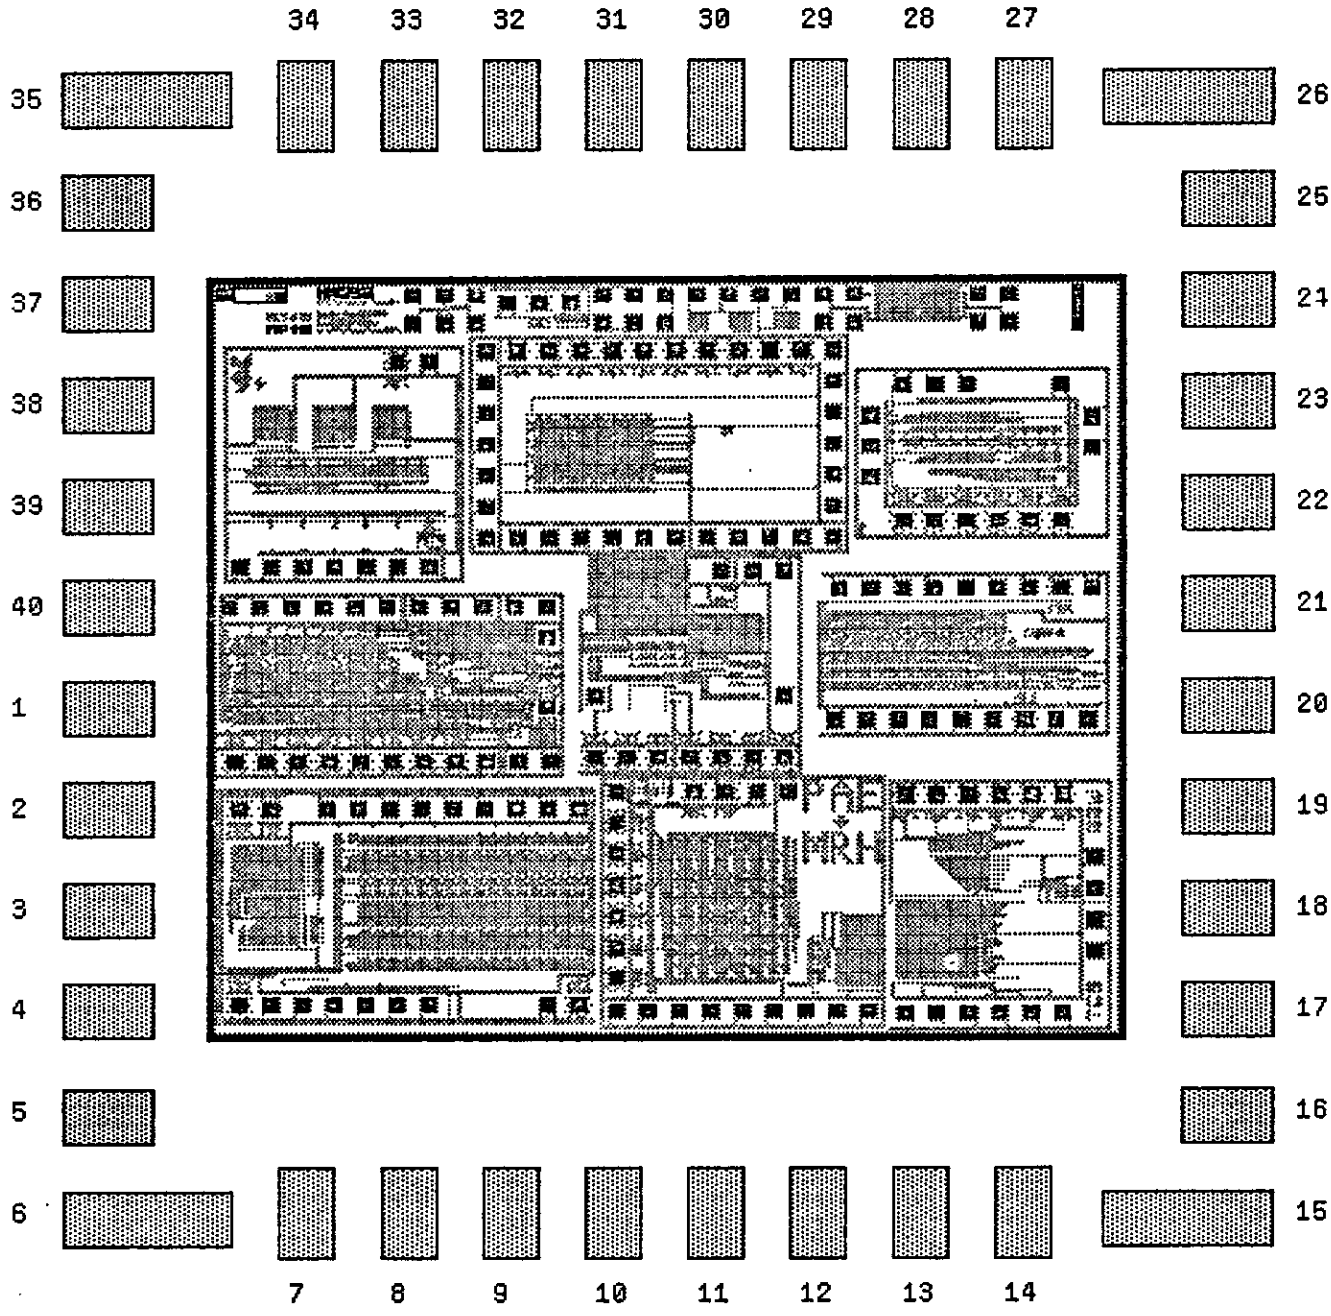
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 BJ

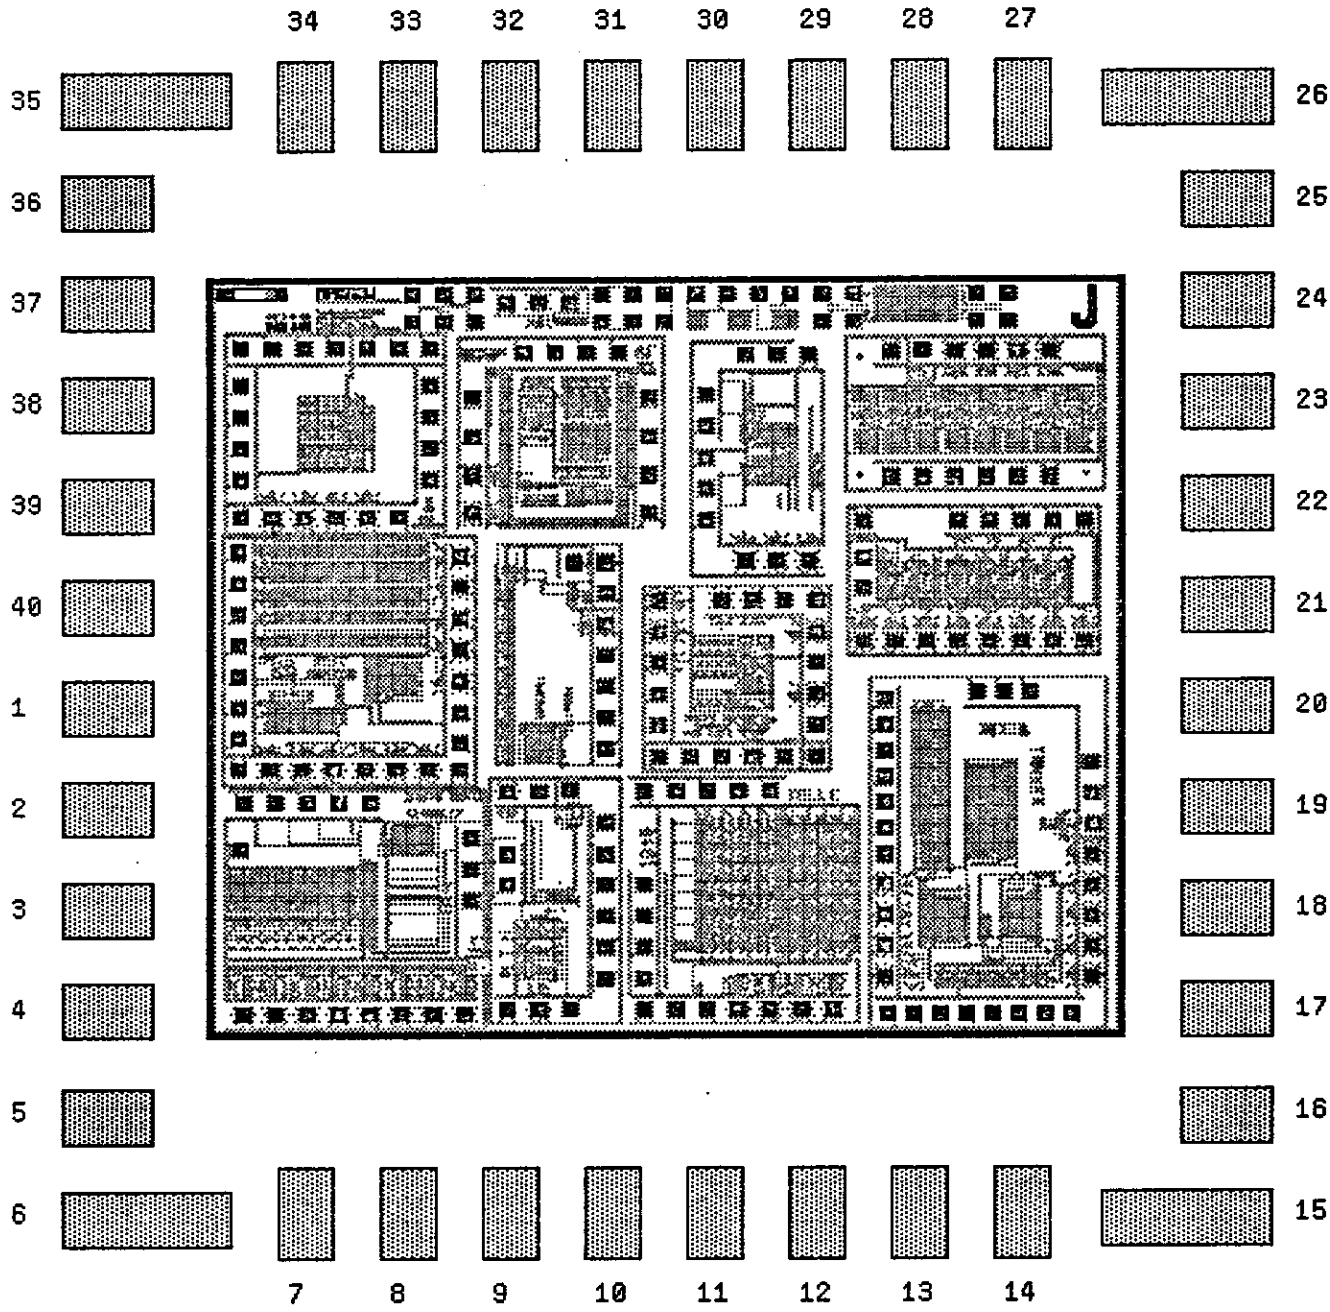
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 BK

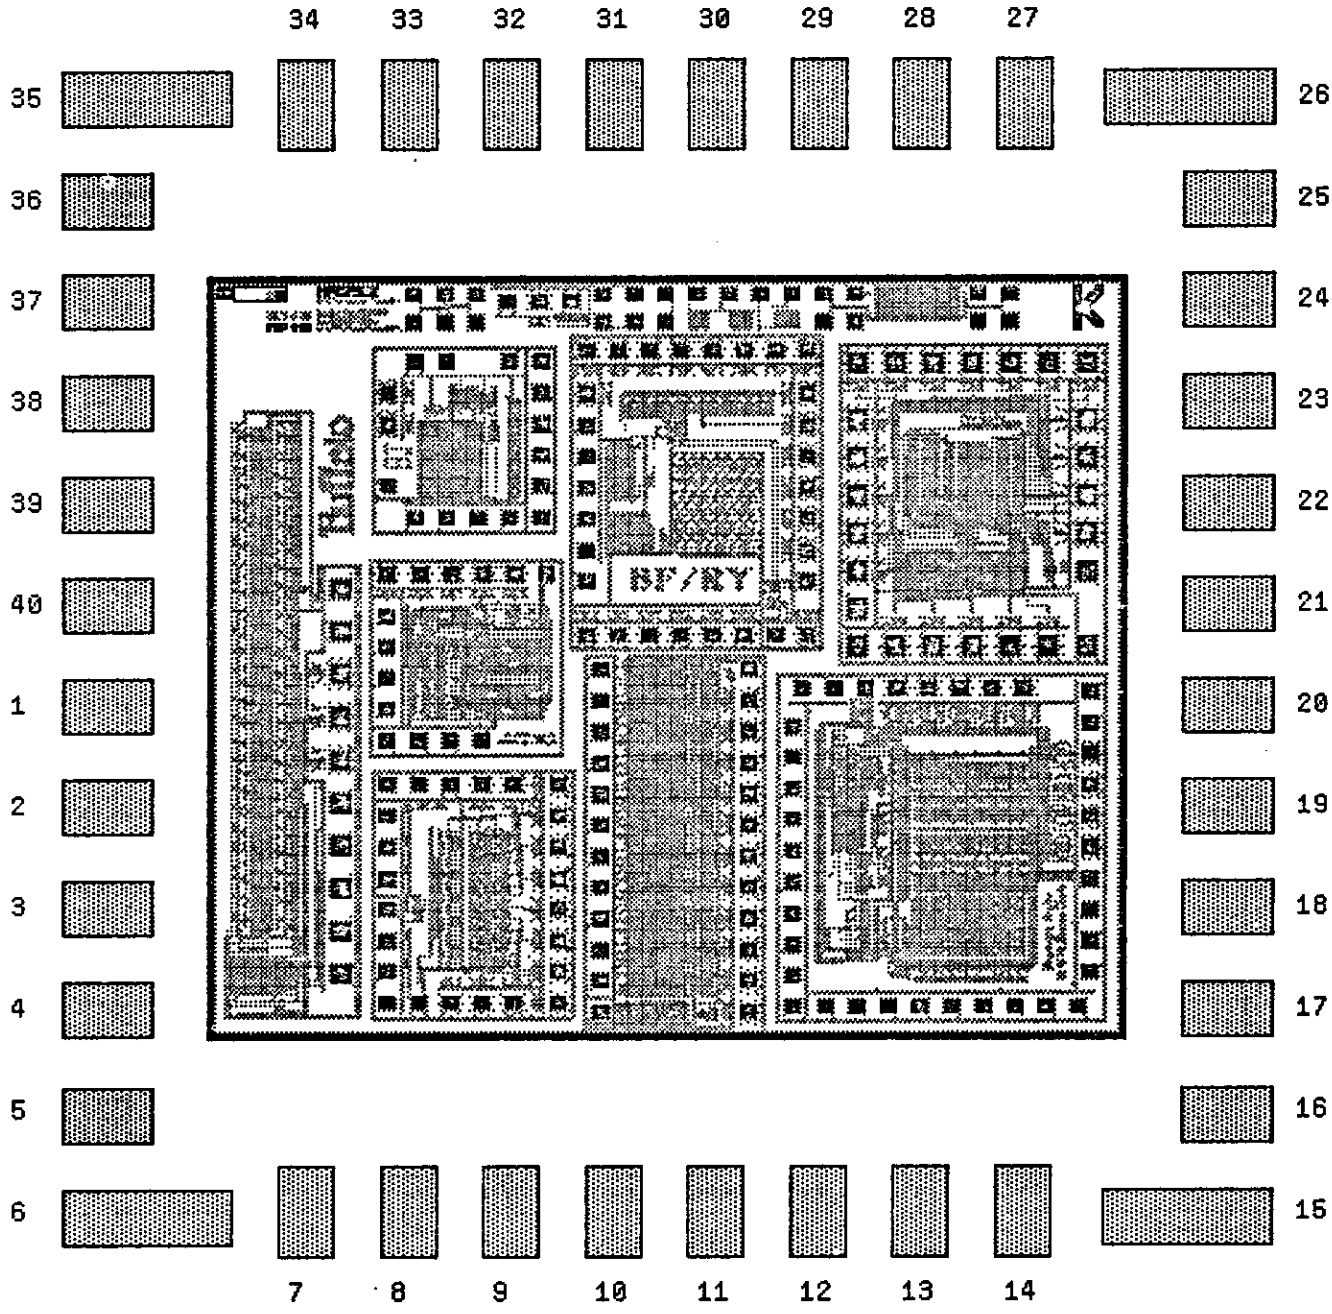
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 BL

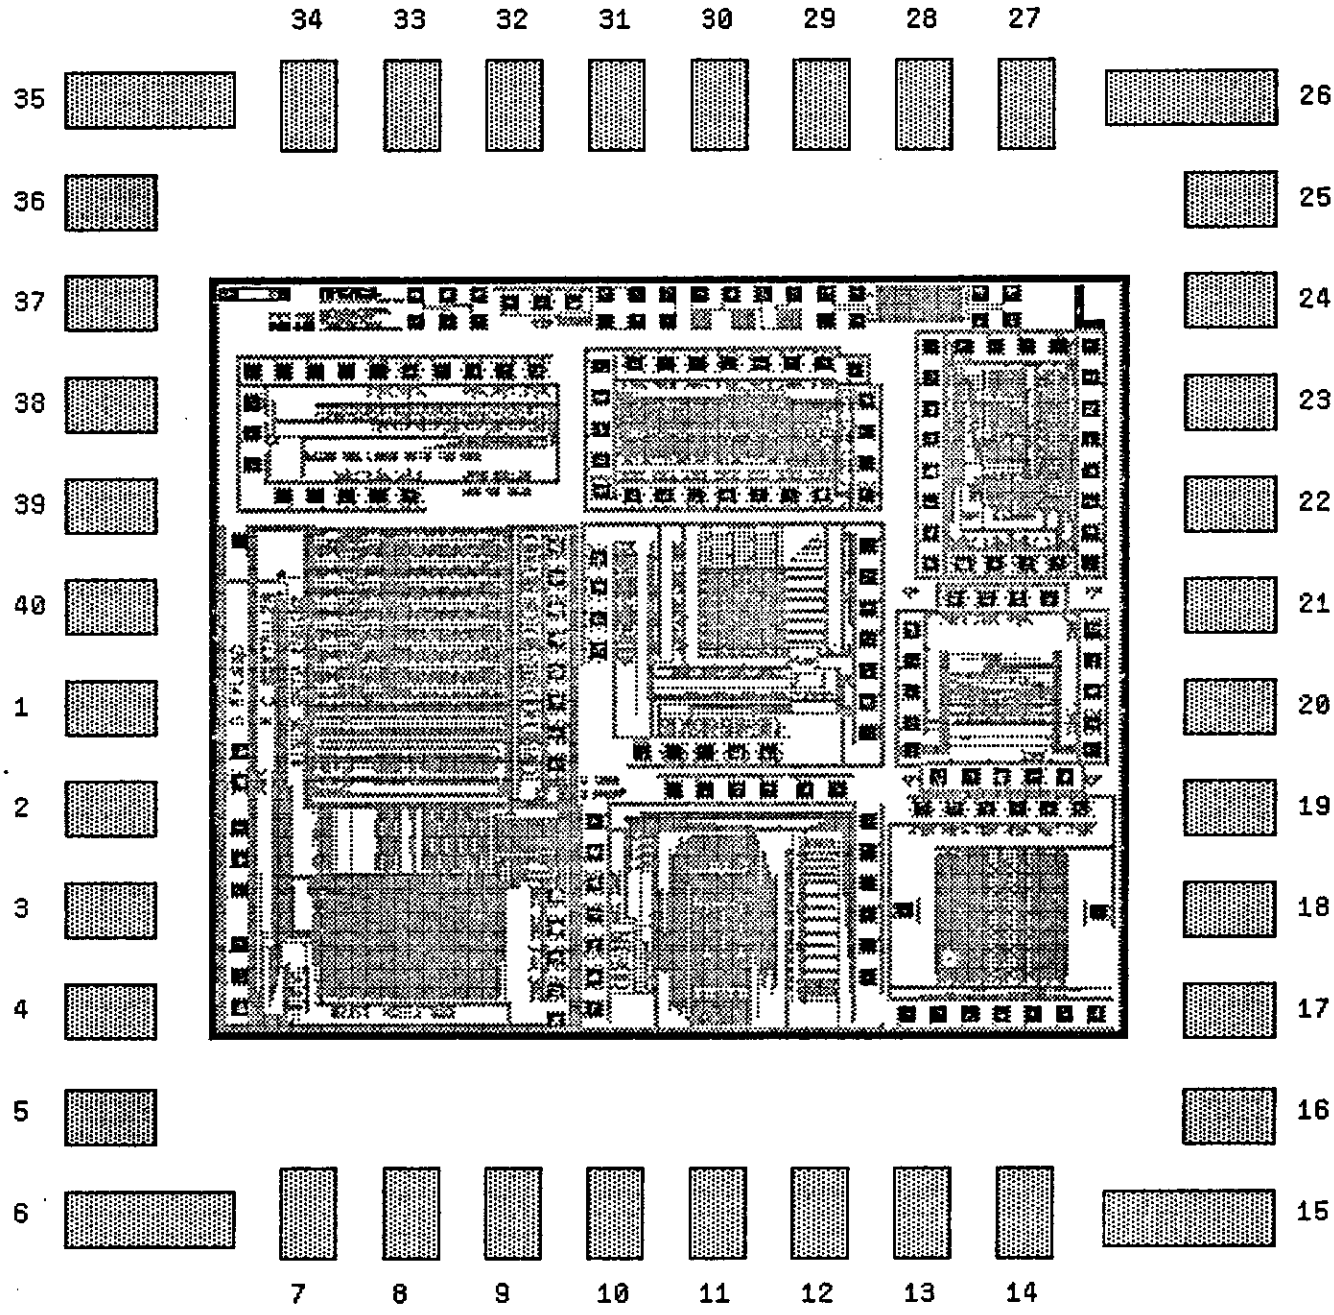
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 8M

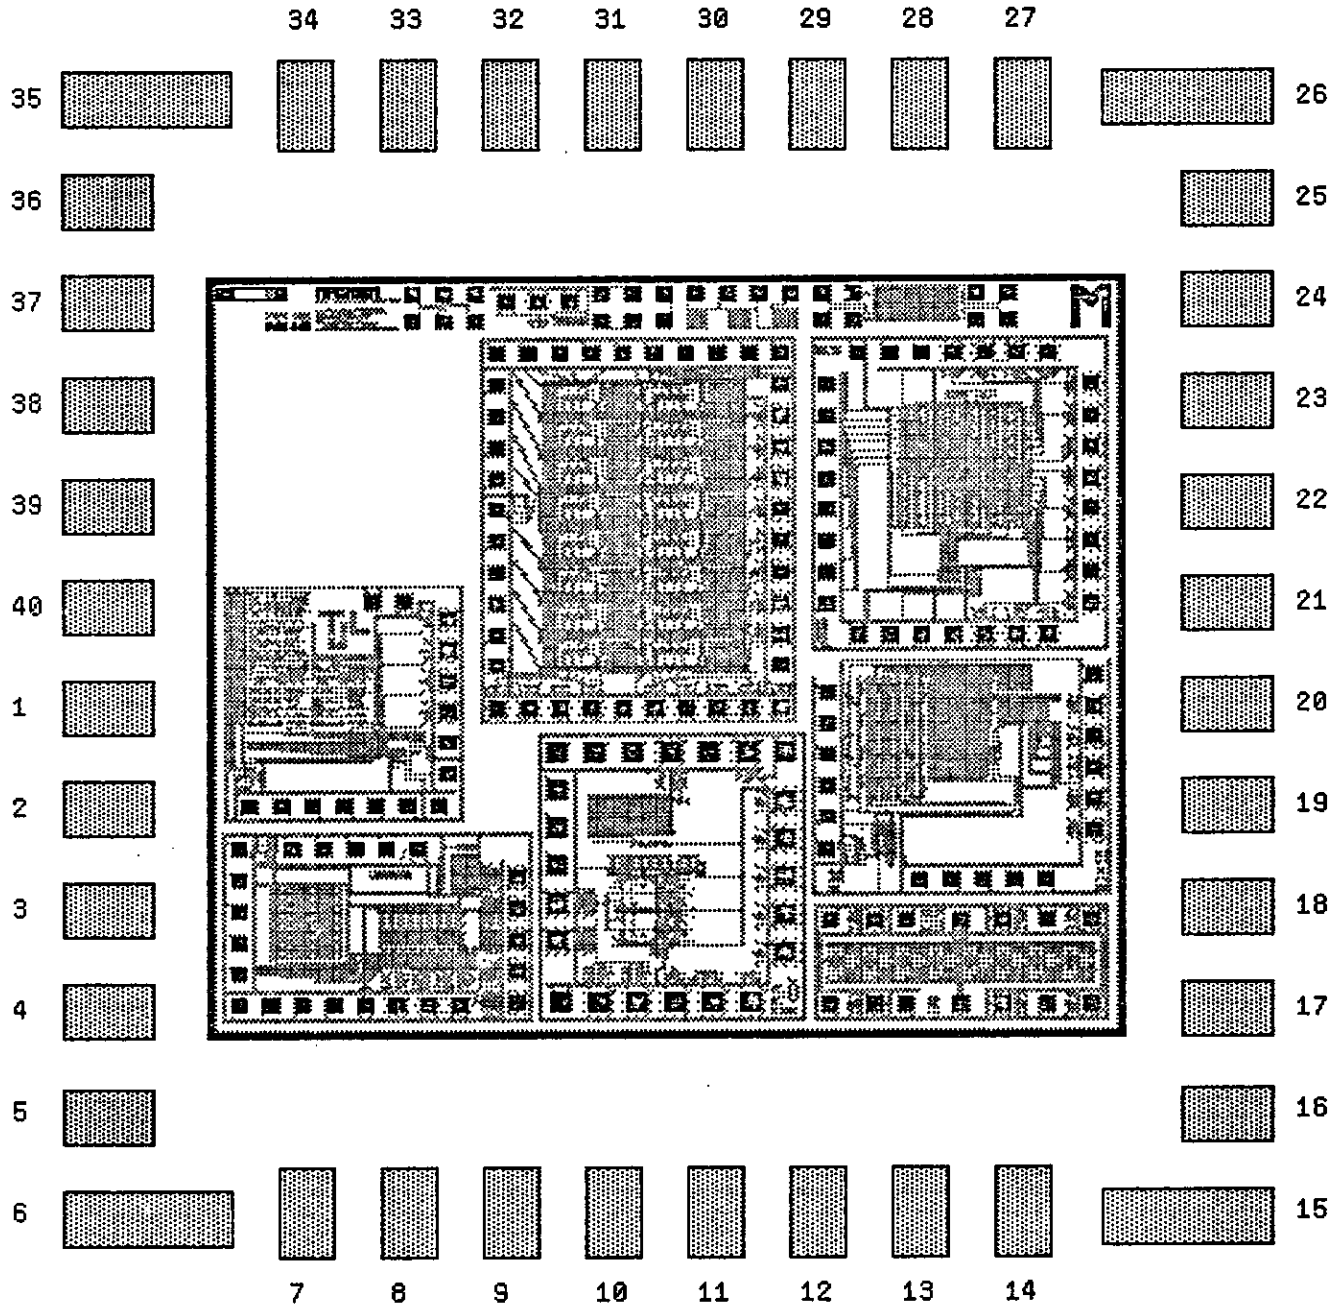
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

MPC79 BN

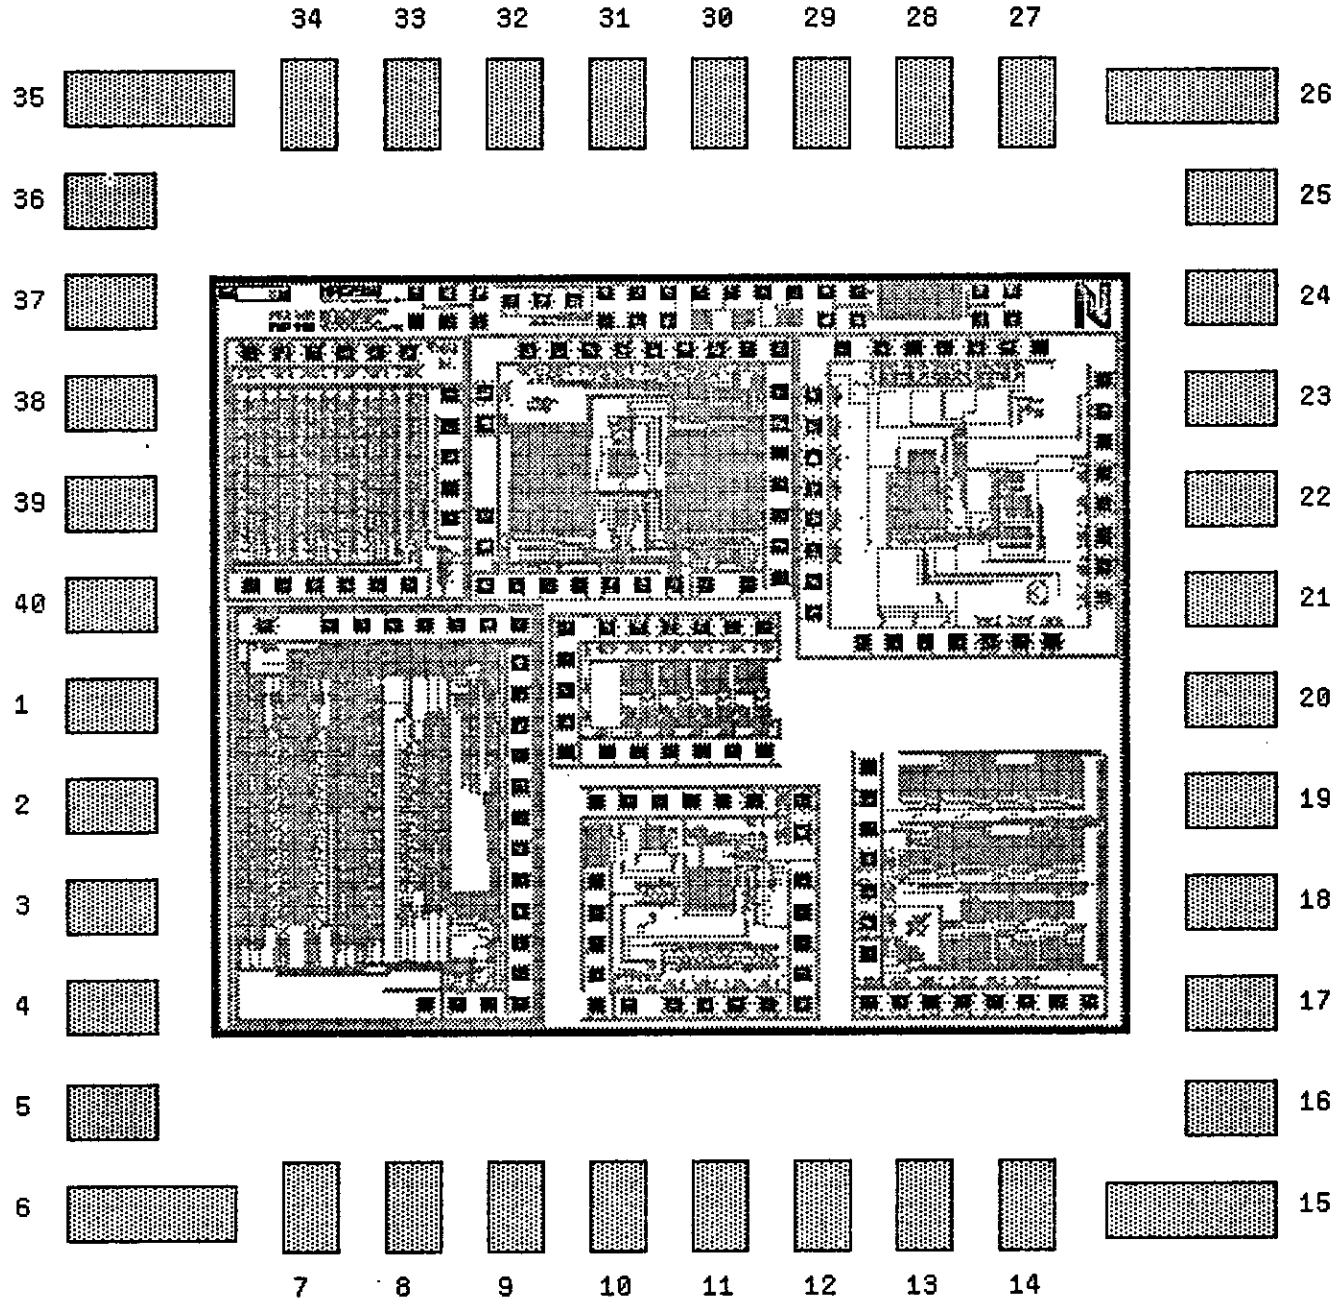
34
33
32
31
30
29
28
27
35
26
36
25
37
24
38
23
39
22
40
21
1
20
2
19
3
18
4
17
5
16
6
15
7
8
9
10
11
12
13
14

7. Further Plans and Request for Feedback

This section presents our plans for further documentation and publication concerning MPC79, and describes how designers can participate by providing us with feedback about their projects. We encourage designers to test their projects, and to report on the results. The designers of interesting projects are encouraged to seek further publication of their work. Also covered are our plans for further packaging and distribution of chips, and for distributing photographs of the chips.

Feedback about Testing

We are very interested in learning the results of the functional testing of projects. We will highlight and strongly feature in our upcoming publications the work of those designers who provide us with the results of the testing of their projects. A questionnaire and a postage-free business-reply envelope is being distributed with this document; these can be used for submitting early test results and for providing us with other important feedback. Any test results that we receive by 31 January 1980 are sure to be mentioned in our upcoming publications.

On Friday, 28 December, shortly after the type-B wafers came off the line at HP-ICPL, the first project was tested and proved to function completely correctly (this was Jim Clark's "Hourglass" project, described in the attached short-form project report). We now know that the fabrication process has been successful, and that there are no systematic errors in the overall MPC79 implementation effort. And so, the time has finally arrived to find out if *your* project works! We're looking forward to hearing from you and learning the results.

The MPC79 Proceedings

We are compiling a proceedings documenting the MPC79 effort, to be published as a Xerox PARC/SSL Report entitled: *Proceedings of the MPC79 Multi-University Multiproject Chip Set Project*, edited by A.Bell, L.Conway, R.Lyon, M.Newell. This proceedings will include (i) photos of all the different die types, (ii) information from the MPC79 ARPANET message traffic, (iii) material from this implementation documentation, (iv) information about the MPC Implementation System, and (v) feedback from the universities. *We especially hope to include a collection of short-form project reports describing a number of the MPC79 projects.*

The MPC79 Proceedings will be printed in large quantities and distributed widely in the universities and in industry. For any of you designers who'd like some visibility, now's your chance! Be aware that there is an intense and growing demand for systems folks who know all about VLSI architecture and design. This demand provides students just finishing degrees with many exciting employment opportunities in industry, provides great opportunities for those interested in teaching, and provides university faculty and research staff members with many opportunities for industrial consulting.

Designers who'd like to see their projects featured in the upcoming MPC79 Proceedings should do the following: (i) Fill out the attached questionnaire, indicating that you plan to submit a report. (ii) Return the questionnaire to us by 31 January 1980, so that we can plan to include your report. (iii) Prepare your report: This can be a short summary of your class project report. Try to provide a hierarchical view of your project, both in words and pictures, but in a compact amount of space. Include some information about yourself, and about the tools used to create your design. Two or three pages of text, followed by two or three pages of diagrams should be adequate for all but the largest projects. See the sample short-form project report appended to this section, to get an idea of what we're looking for. (iv) Send the report to us by 29 February 1980.

Other Places to Publish Your Work

Most of the designers participating in MPC79 are university students, working on projects as part of their VLSI design courses. It might not occur to these students that, by completing a novel VLSI design project, they may have done something worth reporting on. It appears to us that many of the projects have highly innovative architectures and novel design features, and that a few students have undoubtedly done important, original work (of course, there are also some highly visible exceptions!). Thus, it may prove easier than you think to have your work published. We suggest that all students who have successfully completed interesting projects seek publication of their work in appropriate journals, proceedings, or magazines. If you aren't sure how to go about this, talk it over with your instructor or faculty advisor.

Some good places to publish are as follows: Major projects having innovative architectures, especially those that tested successfully, should seek publication as soon as possible as reviewed papers in the appropriate IEEE or ACM technical journals. An alternative to this would be to submit papers to any of the upcoming VLSI design conferences to be held at the major universities. Many projects involved the use of novel design aids; papers on these could be submitted to the annual IEEE/ACM Design Automation Conference. Many of the projects could lead to excellent articles in Electronics, Computer Design, Electronic Design, and any of the personal/hobby computer magazines such as Byte. You may also be interested to learn that Doug Fairbairn has started the new magazine LAMBDA, about VLSI design, for the VLSI system designer. He is quite interested in featuring novel MPC79 projects in upcoming issues of LAMBDA. To contact him write to: LAMBDA Magazine, P.O.Box 50503, Palo Alto, CA 94303.

Plans for Packaging More Chips

Since we have only limited packaging capabilities, we are distributing only one packaged chip per project for most of the projects in the shipments going out on 2 January 1980. In some cases, where we know that projects will undergo immediate functional testing, we are distributing two or three packaged chips per project.

There will be a need for further packaging. Probably the largest demand for this additional packaging will be for projects having multiple designers, so each designer can have a copy. We are shipping a moderate number of unpackaged chips to each coordinator in the shipment of 2 January, and coordinators may be able to make local arrangements for the packaging of these additional

chips. We may do some additional packaging here at PARC, and we may also be able to arrange for further packaging to be done elsewhere.

So, if you'd like to get hold of more packaged chips, talk to your lab coordinators. They may be able to make the local arrangements, or they may relay your request on to us. We'll expedite the packaging of additional chips for those who plan to do some serious functional testing, especially for the larger projects (if your project has a very large area, you may need to test several copies of it, in order to separate possible design errors from possible fabrication defects).

Plans for Making Chip Photographs

In the near future, a commercial photographer will be producing 8" x 10" photographs, both in color and in B/W, of each of the MPC79 die-types. We'll distribute copies of these photographs to each university, along with information on how to order more copies from the photographer.

We may arrange to leave a set of wafers and wafer-maps with the photographer, and try to set up arrangements so that participants can have photographs taken of their *individual* projects, ordering them by ProjectID. Many folks will want to obtain photographs of their individual projects to use with papers they are publishing, to have as souvenirs, or to use as frontispieces for their resumes! Also, some schools may want to compile sets of photos of their successful projects. If we're successful in making the arrangements for ordering individual project photos, we will announce them at the time we return the first set of 8" x 10" overall chip photographs.

Our Plans for Reports on VLSI Implementation and on the MPC Implementation System

During the Spring of 1980, we'll be writing several Xerox PARC/SSL Reports concerning VLSI implementation that may be of interest to those who wonder about what went on "behind the scenes" to pull off the implementation of MPC79. The overall subject of VLSI implementation will be discussed in *The Implementation of VLSI Systems*, by L.Conway, A.Bell, and M.Newell. A detailed description of the internal structure and operational use of the MPC implementation system software will be given in *The MPC Implementation System*, by A.Bell and M.Newell. As soon as these reports are printed, we'll send some copies out to the EE/CS departments at the participating universities.

Conclusion

Good luck on your functional testing. Be sure to send in those test results and project reports! MPC79 has been very exciting and a lot of fun for us. We hope the same is true for you!

Happy New Year!

Lynn Conway, Alan Bell, Martin Newell, Richard Lyon, Richard Pasco
1 January 1980

An nMOS System Clock

by
James H. Clark
Computer Systems Laboratory
Stanford University

Introduction.

This is one of two projects by the author that are included in MPC79. These two projects comprise the principal parts of a Geometry Engine for doing floating-point operations on 4-component vectors. The final chip will consist of one copy of the "Hourglass" project described here (BK-5, Clark2SU) and four function units, each one of which is two copies of the "Geometry Engine" project in MPC79 (BK-8, ClarkSU), with appropriate numbers of the principal bit-slice.

Each of the function units of the final chip has three registers, an ALU and a stack. The units can be microprogrammed for a variety of different functions. Specific microprograms currently being implemented perform four of the computation-intensive tasks common to much of computer graphics: 4x4 matrix transformations; clipping and scaling of polygons, lines and characters; and 4x4 determinants. In one parallel/pipeline organization, a set of 12 identical chips transforms four-component vectors, clips to six planes, and scales three coordinates to the system of the destination display at the rate of about 3500 lines (or 900 four-sided polygons) every 1/30 second, assuming $\tau = .2$ nsec. Some interesting characteristics of the final system are: 1.) a large number of functions are accomplished with a very simple, homogeneous organization, 2.) each chip is locally a synchronous system that communicates asynchronously with other chips, as described in detail below, and 3.) using 1 micron line widths, it will be possible to fabricate the 12-chip system on a single chip, with a factor of 4 performance improvement over the figures given above.

Because the organization described above requires that each chip communicate reliably with other chips, the author chose a self-timed, 4-cycle communication protocol that avoids synchronization failure, as described in Mead & Conway by Chuck Seitz. The clock circuit "Chuck's Hourglass" described here is an adaptation of the pipeline clock given by Seitz that generates two-phase, non-overlapping clock signals.

Description.

The main circuit for the clock is shown in Figure 3. The oval shapes shown in this diagram are asymmetric delay lines. These delay lines propagate high-going levels with a delay governed by the inverter pair delay (plus capacitive loading); low-going levels propagate very quickly [cf. Ch. 7, Mead & Conway]. The plan view of the layout for the chip is shown in Figure 1.

This clock has a very long Phi1 and a relatively short Phi2. By convention, throughout the Geometry Engine, all registers are loaded and ALU's are precharged on Phi2, and computation and bus precharging take place on Phi1. Registers are also refreshed on Phi1. This convention can be likened to a standard TTL clock in which the TTL clock pulses occur on Phi2.

To understand the behavior of the clock, refer to Figure 3 and assume that the clock RUN level has been low sufficiently long for the bottom input to the 3-input NOR gate (call it gate 3n) to be at logic level 0. In this state, the output of 3n will be 0, PAUSE is 1, Phi1 is 1 and Phi2 is 0. A short delay after RUN is asserted, the output of this NOR will be a high-going signal presented to the input of delay line d2, and Phi1 will go low. After the delay of d2, a high is presented to both inputs of the Phi2 gate, causing Phi2 to go high. On leaving delay line p2, the inverter reverses the polarity of the transition, and a high level is very rapidly presented at the bottom input of gate 3n, causing its output to go low, thereby lowering Phi2. On the second time around the loop, the transition is low-going and hence fast through d2 and p2, and delay lines d1 and p1 determine the time required for the signal to propagate. If LongTick is asserted before p1a goes high on the second time around, Phi1 is lengthened by the delay difference between taps p1a and p1b.

This clock differs from that given by Seitz in several ways. First, Phi1 and Phi2 are obtained directly from the delay line; hence it is very easy to select the amount of non-overlap between Phi1 and Phi2. A second difference is the Pause signal. This signal is asserted when the clock has completed a cycle, i.e. the end of Phi1, but is not running. Its use is peculiar to this particular application.

To avoid synchronization failure, the clock incorporates two sets of 4-cycle communication circuits. One set is for normal pipeline communication, similar to that given by Seitz, and the other set is peculiar to the particular way that this clock is used in the present context. Figure 2 illustrates the circuits that handle the 4-cycle handshaking, and their effect on the RUN signal is shown in Figure 4. Incoming requests(ReqIn) start the clock if it is stopped waiting for an input request(LoadInput is asserted). Likewise, the clock stops if the control asserts LoadOutput and the previous output request has not been acknowledged. The behavior of these two blocks is precisely that given by Seitz, but the circuit implemented here avoids his circuit in which "ground" is supplied to an inverter by the output of another inverter, thereby avoiding unusual pull-up/pull-down ratios.

One feature not shown directly but implied by the pad labeled "Fast" in Figure 1 is a level of multiplexing that selects between two global clock periods. By not asserting Fast, everything is lengthened. The intent here was to provide a way that the clock could be slowed down by 30% over the normal speed to allow testing the circuit being driven by it without concern for speed problems.

Test Results.

The clock chip was bonded and tested on 12/28/79. All worked as planned except that the period was approximately a factor of two longer than expected. This is presumably due to having neglected effects of stray capacitances in the delay lines. It is easily compensated for by eliminating some of the capacitive loading in the delay-line cell. The measured periods in nanoseconds were:

LongTick	Fast	d2	Phi2	d1	Phi1
0	1	60	137	50	430
0	0	60	150	50	560
1	1	60	137	50	590
1	0	60	150	50	840

For the purposes of the test chip, one set of ReqAck blocks was disconnected. The other set of blocks were connected as follows: the ReqOut of one block was connected to the ReqIn and LoadIn of another; likewise, the AckOut was connected to AckIn. LoadOut was set to be the complement of the ReqOut. This connection caused a ReqOut to be generated every other cycle of the clock; the corresponding AckIn was generated on the alternate cycles. In this way the clock sent and acknowledged its own requests.

The project layouts were done using the ICARUS interactive circuit layout program. The total design and layout time for the clock was approximately two weeks. The Geometry Engine design time was three months, and layout time was approximately three weeks. A very crude design rule checker was used to check parts of the circuit. Helpful assistance was given by Dick Lyon and Martin Newell, of the LSI Systems group, Xerox PARC/SSL.

Designer: James H. Clark
Position: Associate Professor
Address: Computer Systems Laboratory
Stanford University
Stanford, California 94305
Telephone: (415) 497-1414

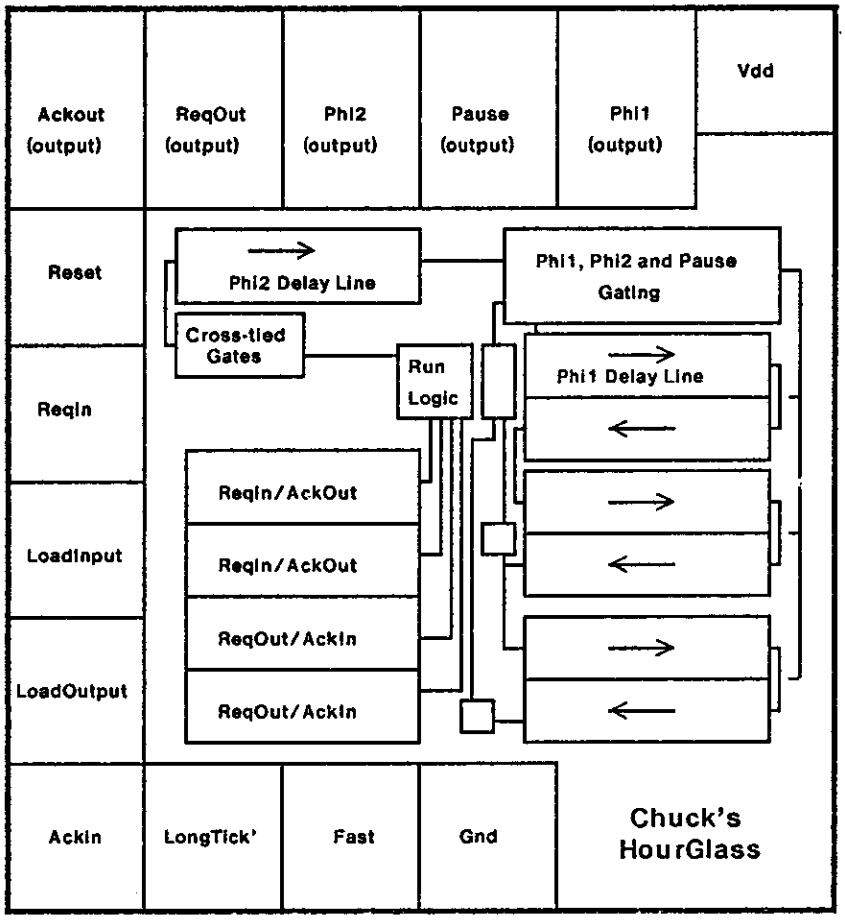

Figure 1. Clock Chip Floorplan, MPC79-BK.

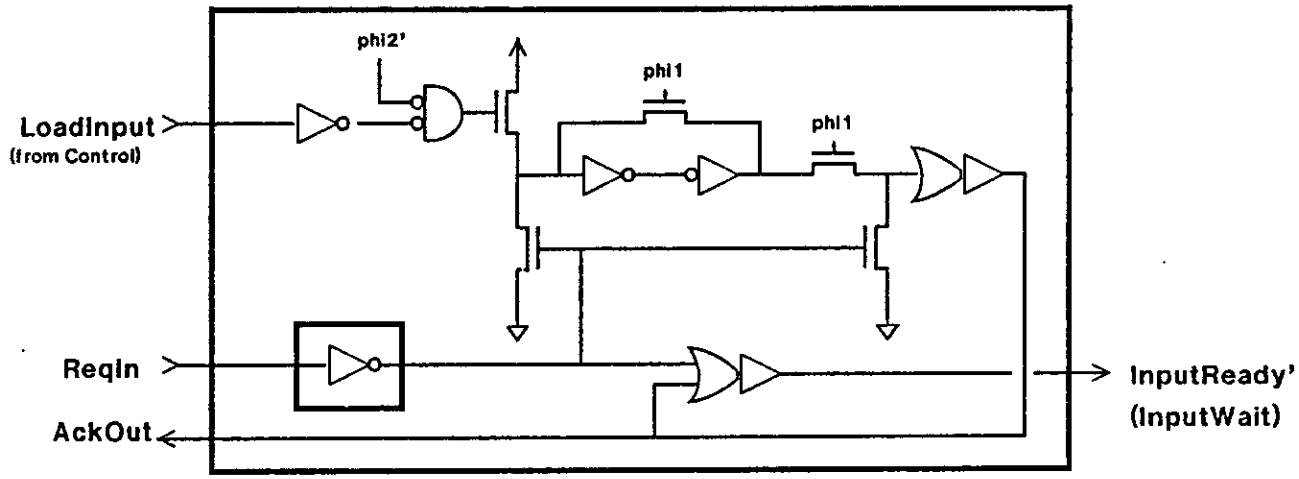

Figure 2. ReqIn/AckOut Block. The Circuit for ReqOut/AckIn is identical except that the indicated inverter is replaced with wire and interchange "out" and "in".

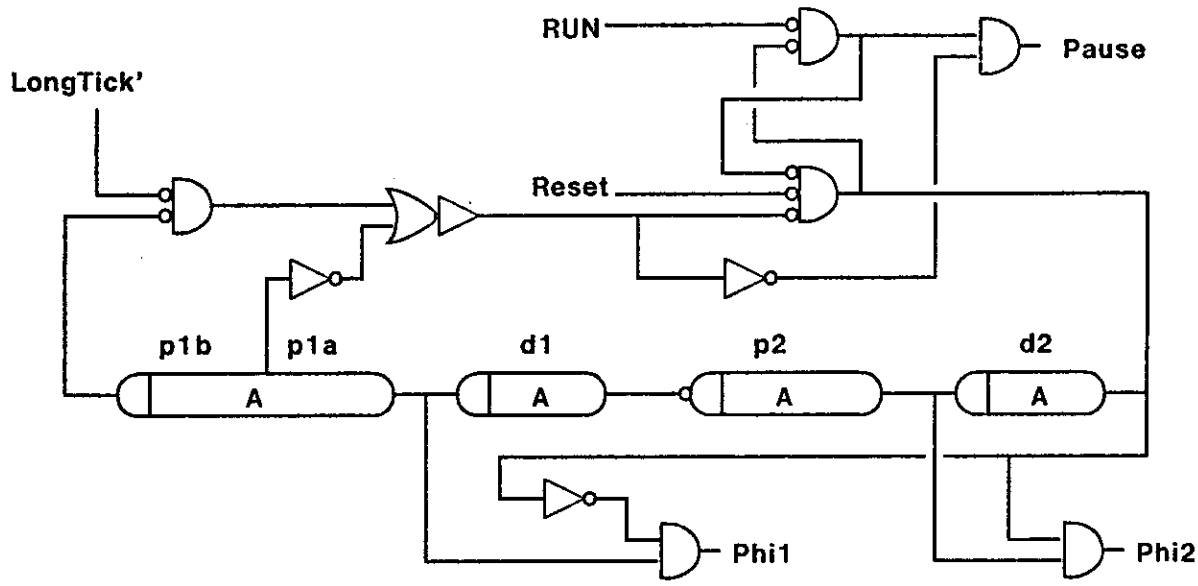

RUN = 0 before end of Phi1 stops clock synchronously with Phi1 high.
RUN = high-going starts clock with Phi1 going immediately low.

d2 = delay from low-going Phi1 to high-going Phi2.
p2 = approximate Phi2 high time.
d1 = delay from low-going Phi2 to high-going Phi1.
p1 = Phi1 high time.

Figure 3. Clock Circuit.

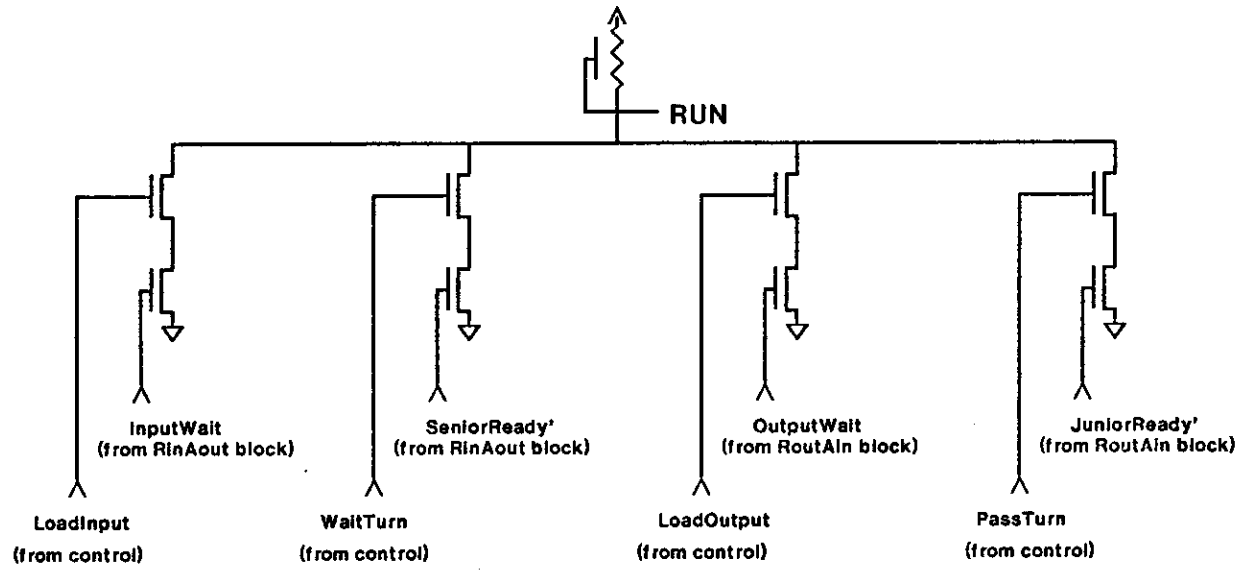

Figure 4. Run Logic.

A Questionnaire for the MPC79 Designers

It would be extremely helpful to the evaluation of the MPC79 prototype implementation system if each participant would take the time to fill in the following questionnaire. We would appreciate both positive and negative feedback - we would like to know what you liked about the service, so that we can be sure to retain those aspects in future implementations, as well as what you did not like together with suggestions for improvements, so that we can try to do better next time.

Please feel free to fill in as much or as little as you desire - all sections, including your name, are optional. If you want to express an opinion about only one topic then fill in just that topic. On the other hand, add additional sheets if you need more space than that provided.

Please mail the questionnaire back to Ms. Lynn Conway, Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304. Addressed, postage-paid envelopes have been included for your convenience. The target date for getting the questionnaires back is January 31, 1980. This is intended to give sufficient time for most people to test their projects.

Your Name: **Project ID:**

University/Department: **Industrial Affiliation (if any):**

Degree sought (if any): **Expected graduation date:**

Have you tested your project?: Yes ☐; No ☐.
If so, did it work?: Fully ☐; Partially ☐; Not at all ☐.
If not, by when do you expect to test it?:
Further testing information:

Will you be submitting a short report of your project,
by February 29, 1980, for inclusion in the MPC79 Proceedings? Yes ☐; No ☐.

Are you planning to report on your project elsewhere?: Yes ☐; No ☐.
If so, in which journal, magazine, or proceedings?:

Did you receive copies of the MPC79 Informational Messages?: Yes ☐; No ☐.
If so, what did you think of their:
Content: Too much ☐; OK ☐; Too little ☐. Clarity: Good ☐; OK ☐; Poor ☐.
Comments:

Implementation Documentation (i.e. that sent along with the chips)
Clarity/conciseness: Good ☐; OK ☐; Poor ☐. Sufficiency: Good ☐; OK ☐; Poor ☐.
Is the information about your project correct in the documentation?: Yes ☐; No ☐.
Comments:

What did you like most about the MPC79 implementation service?:

What did you dislike most about the service? What did you think of working to deadlines?:

Did you find the cell library useful? How could it be improved?:

Any ideas about how to improve the service, to better support the VLSI design courses?:

Other comments (feel free to be verbose, use extra sheets as necessary):

Your Address(es)/ Phone #(s):