A PENTON PUBLICATION U.S. $10.00

TK
7800
E375
Vol. 42
no. 17
1994

ELECTRONIC DESIGN

UNIVERSITY OF CALIFORNIA
LOS ANGELES
SEP 0 1994
ENGINEERING &
MATHEMATICAL
SCIENCES LIBRARY

AUGUST 22, 1994

FOR ENGINEERS AND ENGINEERING MANAGERS—WORLDWIDE

COVER STORY P. 69

STACKED CHIPS PUSH MEMORY VOLUME DENSITY TO NEW HIGHS

- DEBUGGING EMBEDDED SYSTEMS BY USING BOTH TEAMWORK AND TOOLS P. 87
- DUAL-PORT MEMORY CREATES POWERFUL HOST-DSP INTERFACE P. 77
- WHAT'S NEW IN INTERCONNECTIONS P. 129
- PEASE PORRIDGE P. 119

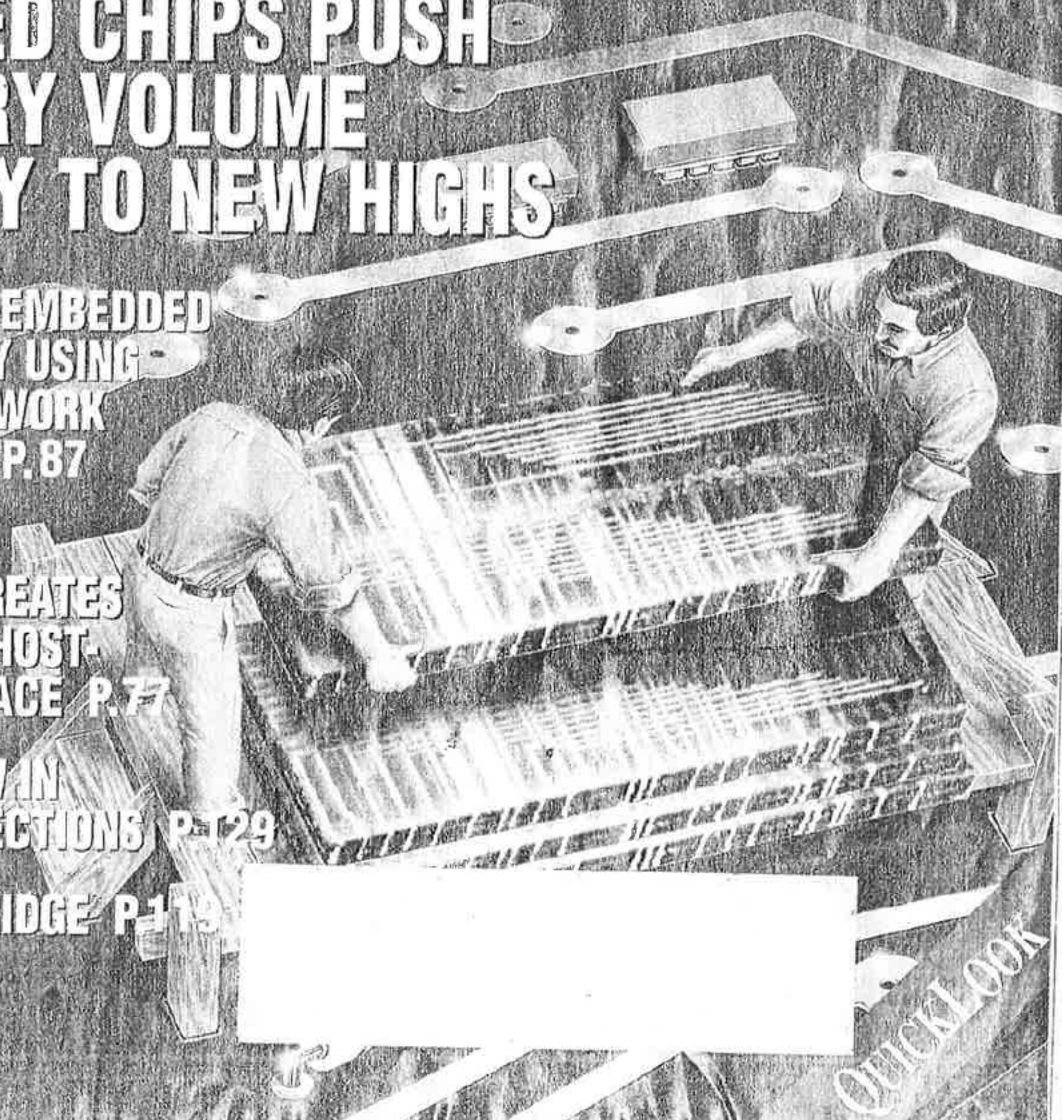

DOCKET ALARM
Find authenticated court documents without watermarks at docketalarm.com.

AUGUST 22, 1994 VOL. 42, NO. 17

ELECTRONIC DESIGN

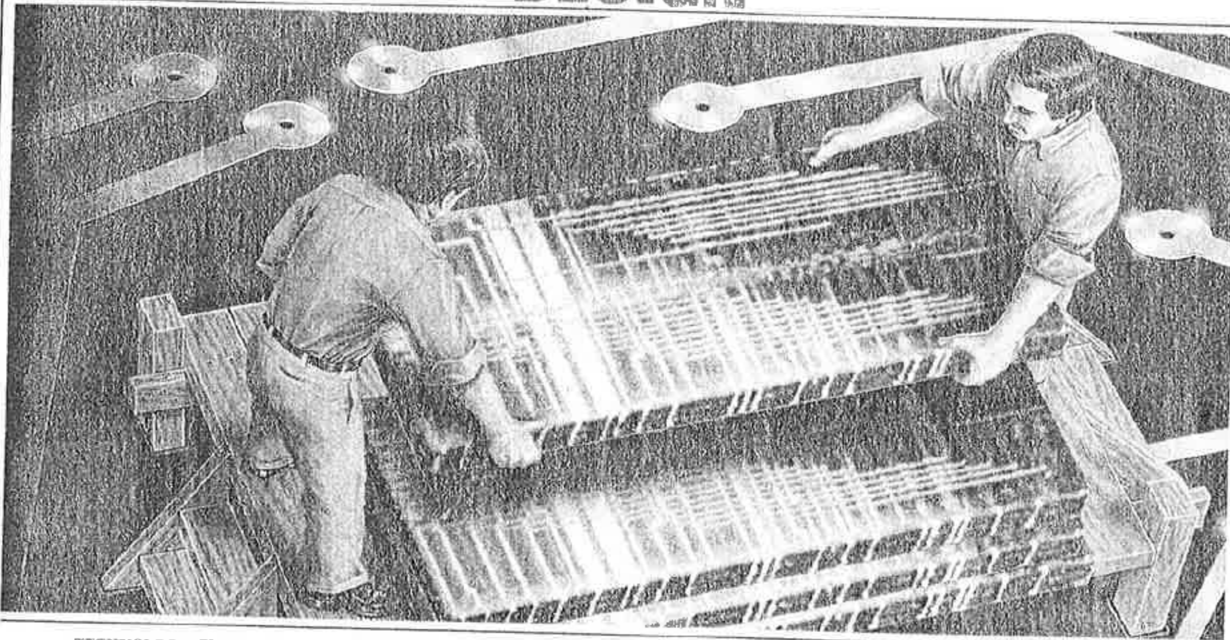

TECHNOLOGY ANALYSIS **HIGH-SPEED MODEMS: THEY'RE NOT JUST FOR DATA ANYMORE . . . 47**
Innovative modem technologies pump data, fax, and voice-mail at blazing speeds. Just how do they get those 114 kbits on that 3-kHz line?

COVER FEATURE **MEMORY-CHIP STACKS SEND DENSITY SKYWARD . . . 69**
Vertical integration of silicon allows packaging of extremely dense system memory in tiny volumes.

DESIGN APPLICATIONS **BUILD A POWERFUL HOST-DSP INTERFACE . . . 77**
Dual-port memory combines the ease and cost of FIFOs with the datapath capability of custom hardware.

DEBUGGING SYSTEMS: GETTING THE PIECES TO WORK TOGETHER . . . 87
Fixing errors in embedded systems calls for teamwork—and tools.

DEBUGGING: ENTERPRISE INSTRUMENTATION . . . 93
Locating real-time bugs in embedded systems is especially tricky while integrating hardware and software.

DEBUGGING FOR REAL TIME: HARD CHOICES . . . 101
Because every tool has an effect on the system, knowing your tools is the best defense.

PRODUCT INNOVATION **INTEGRATED CONTROLLER SUITS LASER PRINTERS . . . 153**
IC includes DRAM/ROM control, DMA, parallel port, interrupt controller, counter-timer, configurable I/O bus, and programmable I/O lines.

ELECTRONIC DESIGN (USPS 172-080; ISSN 0013-4872) is published twice monthly except for 3 issues in May and 3 issues in October by Penton Publishing Inc., 1100 Superior Ave., Cleveland, OH 44114-2543. Paid rates for a one year subscription are as follows: $105 U.S., $185 Canada, $210 Mexico, $255 International. Second-class postage paid at Cleveland, OH, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 611 Route #46 West, Hasbrouck Heights, NJ 07604. Telephone (201) 393-6060. Facsimile (201) 393-0204. Printed in U.S.A. Title registered in U.S. Patent Office. Copyright © 1994 by Penton Publishing Inc. All rights reserved. The contents of this publication may not be reproduced in whole or in part without the consent of the copyright owner. For subscriber change of address and subscription inquiries, call (216) 696-7000. Mail your subscription requests to: Penton Publishing Subscription Lockbox, P.O. Box 96732, Chicago, IL 60693.
POSTMASTER: Please send change of address to ELECTRONIC DESIGN, Penton Publishing Inc., 1100 Superior Ave., Cleveland, OH 44114-2543.

DOCKET ALARM
Find authenticated court documents without watermarks at docketalarm.com.

VERTICAL INTEGRATION OF SILICON ALLOWS PACKAGING OF EXTREMELY DENSE SYSTEM MEMORY IN TINY VOLUMES.

MEMORY-CHIP STACKS SEND DENSITY SKYWARD

DAVID MALINIAK

Many aspects of board design, such as timing, thermal issues, and manufacturability, require innovative thinking. But if there's one aspect of board design that's relatively predictable, it's the layout, which is done in the X and Y dimensions. Chips are, almost without exception, placed side by side in some arrangement, even if the board is double-sided or flexible.

This constant of board design, however, is the great limiter of a board's packaging density. And nowhere is this limitation more glaring than in the layout of system memory, where row upon row of DRAMs sprawl across vast expanses of pc board like some silicon suburban development.

But the push is on to shrink systems, and the traditional layout concepts are becoming an extravagance. Over the years, there have been a number of attempts at exploiting the third dimension in board layout. Chips have been stood on their edges and sandwiched together and there have been attempts at stacking them vertically. But until now, such schemes have been either too expensive or the yields have been poor. In some cases, the density achieved wasn't worth the effort.

With the development of its technology for vertical integration of memory wafers, wafer segments, and individual die, Cubic Memory Inc. (CMI), Scotts Valley, Calif., has shattered the density barrier *(Fig. 1)*. Instead of the 40-to-80-Mbyte/in.^{3} storage densities possible with conventional packaging—using 16-Mbit memories in small-outline J-lead packages (SOJs) or two-sided thin small-outline packages (TSOPs)—CMI claims the ability to achieve densities of a gigabyte or more per cubic inch.

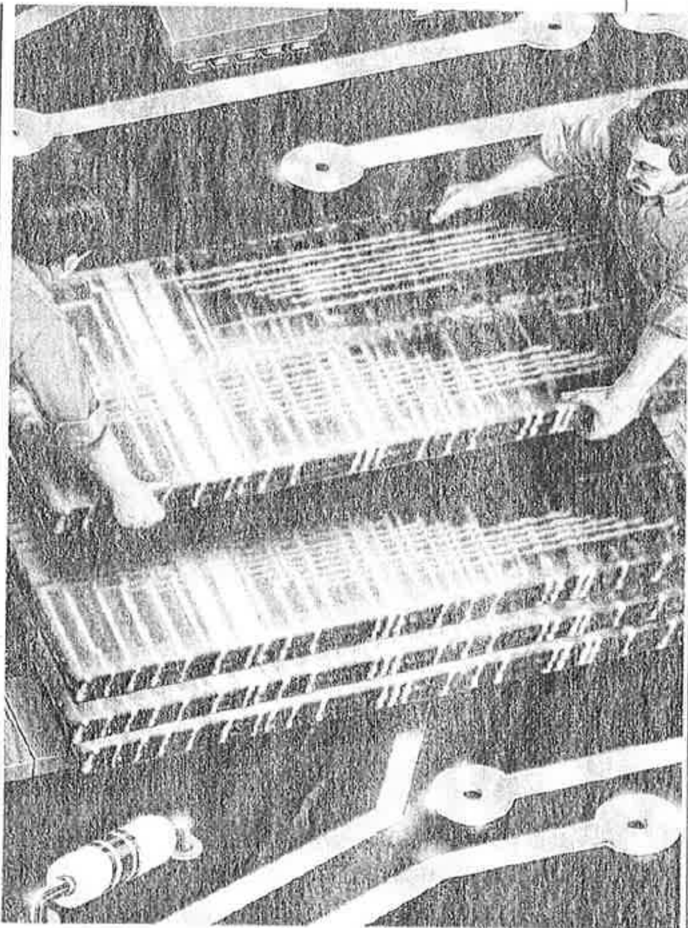

Stacking memory ICs vertically has other benefits besides the density gains. The close physical proximity of the chips significantly reduces the system delay associated with interconnect capacitance and pc-board trace length. Speed is increased while power requirements and operating temperatures are reduced compared with horizontal layouts.

CMI has leveraged its chip-stacking technology into an initial pair of product families.

DOCKET ALARM Find authenticated court documents without watermarks at docketalarm.com.

One takes advantage of the PCMCIA-card format for adding workstation-quantity memory to Pentium-based portables, while the other works within the DRAM SIMM format. Both merely hint at the technology's potential, which is not necessarily limited to stacking of memory alone.

There are actually three separate technologies for vertical integration, each of which is suitable for different applications. The three processes are referred to as the large-hole process, the small-hole process, and the Vertical Interconnect Process (VIP). Common to each process is a vertical silicon interconnect and a patented compliant interconnect scheme.

The first two processes, known as the large- and small-hole processes, use a patented, pyramid-shaped via through the silicon *(Fig. 2)*. The small opening on the top of the pyramid penetrates on the circuit side and the large opening comes through the back side of the silicon, where the interconnect makes contact with a number of circuit elements on the silicon immediately below. This interconnect method takes up no space, as all pins fan out under the silicon stack instead of taking up large areas around the perimeter of the active silicon, as is done with conventional packaging. In addition, this method allows the individual circuit elements to be easily isolated for testing up until the final stack assembly.

The VIP process is a less expensive extension of the small-hole process and does not require the pyramid-shaped via through the silicon. Instead, gold interconnect traces and vertical-interconnect pads are deposited over insulating layers of polyimide. The gold traces provide horizontal interconnection of the die on a single layer or segment, and the gold pads provide a non-oxidizing metal surface for the layer-to-layer interconnections.

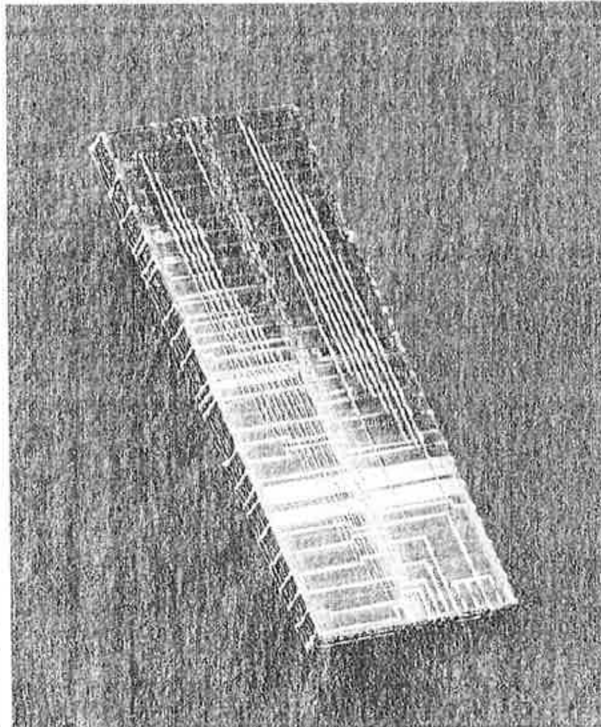

1. INSTEAD OF THE 40-to-80-Mbyte/in.^{3} storage densities possible with conventional packaging (using 16-Mbit memories in SOJs or two-sided TSOPs), Cubic Memory Inc. (CMI) claims the ability to achieve densities of a gigabyte or more per cubic inch.

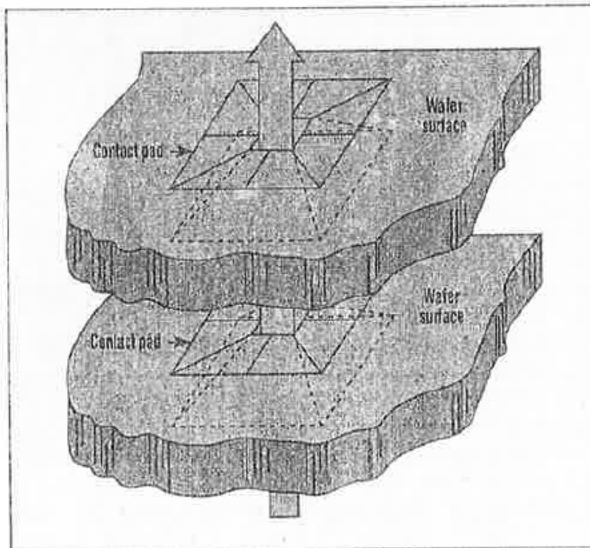

2. BOTH THE small- and large-hole processes use a pyramid-shaped via through the silicon to make the vertical interconnections. The small opening on the top of the pyramid penetrates on the circuit side and the large opening comes through the back side of the silicon, where the interconnect makes contact with a number of circuit elements on the silicon immediately below.

Whole wafers, segments of wafers, or individual die can be stacked. In addition, multiple silicon technologies can be mixed in the same stack *(Fig. 3)*. Other components can be placed on top of the silicon stack as well.

For the small-hole and VIP processes, standard wafers can be used without need for any custom silicon. Individual chips can be selected by the control circuitry. The vertical chip-to-chip distance is variable from 0.003 to 0.010 in., depending on the application and density required.

The vertically integrated stacks are attached directly to pc boards. Tests on whole wafers and wafer segments have shown a 15% reduction in power consumption and a 20°C lower continuous operating temperature when compared to identical circuits operating in conventional plastic packages.

Originally, the large-hole process was developed to allow stacking of complete 6-in. DRAM wafers. The pyramid-shaped holes are filled with a mesh of fine gold-plated wire. The wires form a mushroom shape on the circuit side of the wafer, and the other ends of the wire are compressed and contained in the pyramid base. When the wire-filled base of the pyramid of one wafer is brought into contact with the top of another wafer, the mushroom-shaped "fuzz button" is captured by the base of the pyramid. All of the wires are then contained and a compressive gold-to-gold contact is formed between the circuit elements of both wafers. The result is a very reliable, fully compliant, reworkable interconnect that supports interconnections in both the horizontal and vertical directions.

DOCKET ALARM Find authenticated court documents without watermarks at docketalarm.com.

The "fuzz-button" interconnect does, however, require a custom mask set. And the holes are large by semiconductor standards: 150 μm at the top and 1000 μm at the base. This interconnect technology is still used in some applications.

The small-hole process was developed as a refinement of the large-hole, "fuzz-button" technology to allow interconnects on standard semiconductor-die bonding pads without the need for a specialized structure. In this method, off-the-shelf memory wafers can be used without the need for the semiconductor manufacturer to run a custom mask set.

In this case, the hole-making process is similar to the fuzz-button process, but the dimensions change. The wafers are thinned and the holes are small enough (25 μm on top, 120 μm on bottom) so that the vertical interconnect is accomplished within a bond-pad area and on conventional bond-pad pitches.

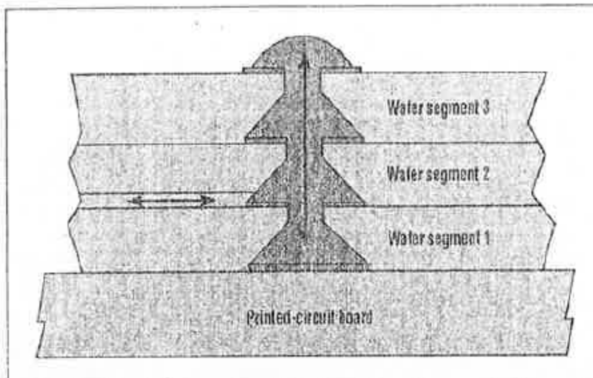

3. WHOLE WAFERS, segments of wafers, or individual die can be stacked. In addition, multiple silicon technologies can be mixed in the same stack. Other components can be placed on top of the silicon stack as well.

The small-hole process incorporates one level of discretionary wiring for every layer of silicon in the stack. Each die also has the necessary control signals brought out so that the control circuitry can address each chip uniquely.

Silver-filled epoxies are used for the compliant conductive material. These proven materials are the same ones that have been used for years by semiconductor manufacturers as a die-attach medium. CMI has developed a proprietary application method for the dispens-

PUTTING MEMORY STACKS TO WORK

Manufacturers of laptop/notebook computers are forever facing the challenge of squeezing desktop performance into an easy-to-transport package. In the RISC-workstation arena, the problem intensifies, because many of the Unix-based applications demand tremendous amounts of memory to execute quickly.

Just such a problem was faced by designers of the SPARCbook III at Tadpole Technology plc, Cambridge, U.K., explains George Grey, the group chief/executive officer. The need for an alternative to the standard DRAM single-in-line memory module (SIMM), which peaks in capacity at 32 Mbytes, led Tadpole to team with Cubic Memory Inc. (CMI) to develop higher-density SIMMs.

By taking advantage of CMI's vertical interconnect process (VIP), designers at Tadpole and CMI defined an extension to the standard DRAM SIMM that will initially provide 64 Mbytes of storage (organized as 16 Mwords by 36 bits, which includes byte-parity bits). Furthermore, CMI expects to double the capacity to 128 Mbytes per SIMM by late 1994, and still keep the SIMM height to just 1 in.

One of the first steps designers had to take to enhance the SIMMs was to expand the address range. They added one more address line and two additional row-address-strobe (RAS) lines by using several pins that were previously no-connect pins on the SIMM. Some of the key issues that designers had to deal with, explains Dave Pedersen, V.P. of engineering at CMI, included the banking architecture and capacitive loading effects in highly-configured systems. Fortunately, because the memory chips are mounted on top of each other, there are no package loading effects. Also, typically, only one bank of memory is turned on at a time, which minimizes thermal problems.

Once the interface was defined, the memory structure was implemented using "stacks" of memory layers. The memory stacks consist of eight-segment layers (each segment can be thought of as a 2-Mword bank of DRAM). The VIP scheme employs pyramid-shaped recesses on the outer edge of the chips. Then, when the chips are stacked, the contact edges are exposed so that conductive epoxy easily fills the recesses and realizes the connections.

Each segment layer in the memory stack consists of a monolithic piece of silicon that contains four 2-Mword-by-8-bit DRAMs, thus forming a 32-bit-wide memory block. Complementing the 32-bit-wide stack is a second stack that provides the four parity bits for each word.

After the stacks are assembled, they are mounted in the SIMM substrate. The substrate actually has a hole the size of a stack cut into it and the stack is then inserted into the substrate. The stack protrudes out of one side of the substrate by the amount of the difference in their thicknesses. In the case of the 128-Mbyte stacks, which stand 0.160 in. tall, they protrude 0.098 in. from the 0.062-in.-thick substrate. That's less than the height of a thin small-outline package. When coated with the sealing epoxy, the mechanical durability of the SIMMs is as good as previous-generation SIMMs that were populated with surface-mounted components.

BY DAVE BURSKY

DOCKET ALARM Find authenticated court documents without watermarks at docketalarm.com.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.