The Scalable Pipeline Architecture behind HP's T300 Color Inkjet Web Press

Lluis Abello; Hewlett-Packard Co; San Diego, CA / USA

Abstract

The HP T300 Color Inkjet Web Press uses two identical print engines capable of printing full variable data at 400 fpm and 1200x600dpi on a 30" wide web. Each print engine prints four colors of ink (CMYK) and Bonding Agent (BA). Each of the five inks employs two print bars in tandem in an arrangement that provides nozzle redundancy. Each of the 10 printbar contains seven 4.25" thermal inkjet printheads with over ten thousand nozzles each. This massive array of nozzles consumes 35 Giga bits of nozzle firing data per second.

The data flow starts with the Digital Front End which provides the interface to the press, processes job ticketing and decomposes incoming PDF files into small chunks that are then concurrently processed by a scalable array of raster image processors (RIPs) built on high-performance blade servers. This solution allows RIPs to be tailored to meet any demand required by a customer's print job.

The raster images are then compressed and buffered before being delivered in print sequence order to the two print engines.

The print engines partition the incoming raster images into bands called slices; each slice covers the width of one Printhead plus some overlap area. Image processing hardware runs in parallel for each slice transforming the compressed CMYK continuous tone raster into real time nozzle firing instructions.

The result is a scalable image processing architecture that extends to wider presses by simply replicating the basic hardware building blocks across the web.

Scalable Printing Architecture

The HP T300 Color Inkjet Web Press uses 4.25-inch thermal inkjet printheads built with HP Scalable Printing Technology.

The printhead shown in Figure 1 consists of five thermal inkjet printhead die placed on a ceramic substrate and a backend assembly that provides mechanical alignment, ink pressure regulation, and electrical interconnections.

The printhead has two columns of 5,280 nozzles, each producing a 4.25-inch swath with 1,200 nozzles per inch. This gives a printing resolution of 1,200 dpi across the web.

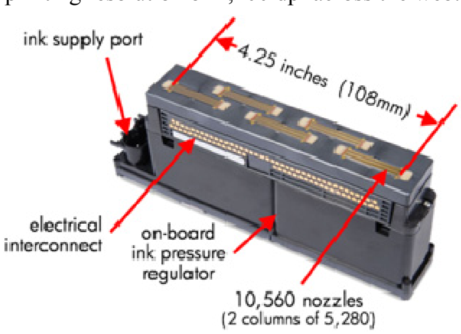

***Figure 1** Scalable Printing Technology*

Printheads are physically arranged into printbars that span the width of the web as shown in Figure 2. Seven printheads are needed to cover a 30" wide web. Printbars are modular and can be added along the web to build different writing system configurations.

A monochrome press can be built with a single printbar, while a color press would have at a minimum 4 printbars one for each of the CMYK primary colors.

The T300 uses two printbars for each color for increased nozzle fault tolerance and load sharing. This allows a 600 dpi pixel in a given dot-row to be printed by any one of eight (8) nozzles, providing the benefits of multiple-pass printing in a single pass.

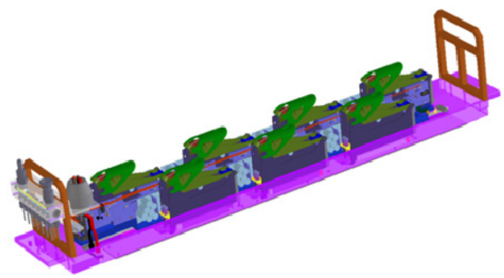

***Figure 2** seven printhead printbar for a 30" wide web.*

Bonding Agent

The HP T300 series Inkjet Web Presses deliver high print quality and durability on uncoated stock by pre-printing a colorless liquid, called Bonding Agent, into those pixels that will receive ink. Bonding Agent is applied by two additional inkjet print bars before colored inks are printed.

The Bonding Agent chemically reacts with the pigment inks to rapidly immobilize pigments at or near the paper surface to control ink spread and penetration, this results in increased optical density and reduced feathering and strike-through. Bonding Agent also improves pigment adhesion to the paper fibers for better print durability.

Nozzle Data Rate

To print at 400 feet per minute (80 inches/s) with a down web printing resolution of 600 dpi, firing data for each of the 10,560 nozzles has to be fed to each printhead 48,000 times per second. This amounts to a data rate of 500 mega bits per second per printhead.

Each of the two print engines needed for duplex printing employs 70 printheads arranged in 10 printbars as shown in Figure 3. Such a writing system contains 739,200 nozzles that consume a total of 35 Giga bits per second per print side.

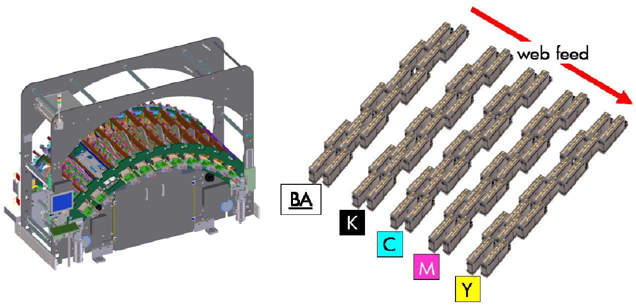

Figure 3 Scalable Print Engine Architecture of the T300 Inkjet Web Press

Scalable RIP Architecture

The data pipeline challenge starts with the Digital Front End which provides the user interface to the press. The DFE processes job ticketing and decomposes incoming PDF files into small chunks that are then concurrently processed by a scalable array of raster image processors (RIPs) built on high-performance blade servers. This solution allows RIPs to be tailored to meet any demand required by a customer's print job.

A monochrome book printing application can use as few as 8 RIPs, while a full-color 100% variable Direct Mail application may use over 100 RIPs to produce 2600 A size color pages per minute.

The RIPs apply the imposition instructions to generate web wide raster images called frames. RIPs also perform color management and ink limiting by applying an ICC profile that is selected on a per job basis. The result is a 600 dpi continuous tone raster frame that is then compressed using a visually lossless compression algorithm leveraged from the Indigo presses. Typical jobs achieve a compression ratio of at least 15:1.

The Press Interface Adapter and Frame Broker interface the Digital Front End and RIPs to the print engine by buffering frames from the RIPs and delivering them in print sequence order. Buffering allows completing the rendering of frames in any order thus maximizing RIP performance. Also when multiple copies of a job are to be printed, the job needs to be RIPed only once.

Standard 10 Gbit Ethernet links are used to transfer the compressed raster between the different sub systems.

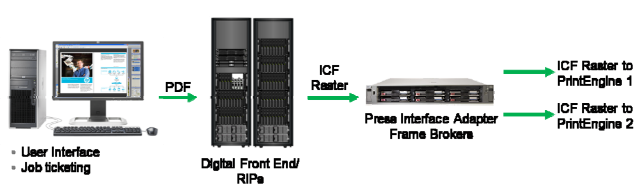

Figure 4 Scalable RIP architecture

Scalable Processing Architecture

The design philosophy behind the scalable printing architecture based on the scheme of spanning the web with overlapping printheads extends also to the scalable image processing hardware.

As shown in Figure 5 an image wider than a single printhead is partitioned into overlapping print swaths, called slices [1][2].

Figure 5 Slicing an image

Printed with tandem printbars for each color (and bonding agent), each slice is 5½ die wide (4.7 inches) to allow overlap between the printbars. Neighboring printheads spanning the web print in an overlap zone for seamless stitching. Modular image processing hardware and firmware components run in parallel for each slice. These components are then replicated across the web to build wider presses.

Front End Pipeline

The image pipeline is split into two distinct steps. The first step is known as the front end pipeline and is performed by the Engine PCA. Here a slice of compressed raster is decompressed, linearized and halftoned into 600 dpi, 3-level pixels for each of the CMYK color planes. If bonding agent is used, the CMYK planes are then merged together to generate the bonding agent plane. The halftone data for each slice and color plane is then transferred to the backend pipeline.

As there are no data dependencies among slices and the compression algorithm allows the raster rows of the frame to be indexed in the compressed domain, the slices can be processed concurrently by each of the engine PCAs.

When printing at 400fpm, each 4.7 inch slice consumes 137 million pixels per second for each of the colors. Each Engine PCA actually processes 4 colors plus bonding agent for 2 slices or a total of 1.37 billion pixels per second.

The Engine PCA has a PCI express interface and form factor. To build a 30" web press 4 such PCAs are required per print side. The PCAs are arranged inside an HP Proliant server.

Back End Pipeline

The second step known as the back end pipeline is performed by the Color Plane Processor PCA. In this step nozzle firing instructions are generated in real time for each of the printhead dice.

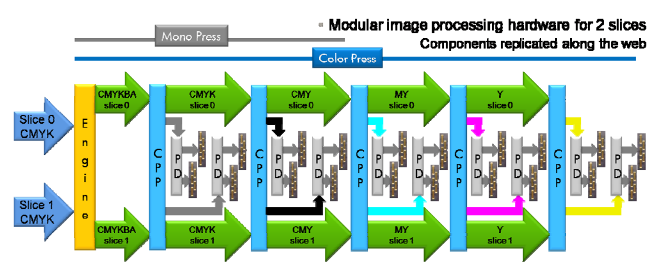

Figure 6 Modular processing hardware for 2 slices and 5 colors

This component is also modular, being replicated down the web to support presses with more printbars. The hardware configuration for a a 2-slice, CMYK and Bonding Agent writing system producing a print swath for two slices (about 8.5 inches wide) is shown in Figure 6. Each CPP processes one color plane and sends the remaining planes to the next CPP.

A 600x600 per inch halftone pixel is formed from two subpixels, each 1/1200" across the web and 1/600" along the web. Printheads can place a dot in either or both subpixels. Halftone level 0 is an empty pixel. Level 1 uses one drop of ink by printing a dot in either subpixel. Level 2 uses two drops of ink, one in each subpixel.

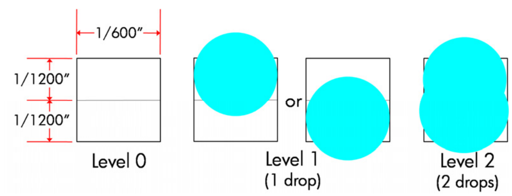

Figure 7 *Halftone Pixels 3 levels per color*

For each printhead die, the back end pipeline reads the corresponding portion of the halftone image that is to be printed by that set of nozzles. The physical position of the die is taken into account to align the dots on paper relative to the other printheads with a 1,200th of inch accuracy. Then the nozzle masking process determines which of the 8 nozzle(s) will actually be used to print each pixel. To perform this, a per nozzle 2 bit mask is combined with the 2 bit halftone pixel and a decision of whether to fire a drop is made. For a given halftone pixel, this process is repeated 8 times, once for each of the nozzles that can print it.

Figure 8 is a schematic representation of the tandem printhead arrangement. Printheads 1 and 2 are stationary and the web moves under them. Nozzles in each column (e.g., "a", "c", "e", etc.) are spaced 1/1200-inch apart to print at 1,200 dpi across the web.

Nozzles "a" and "b" on Printhead 1 print in the same dot-row.

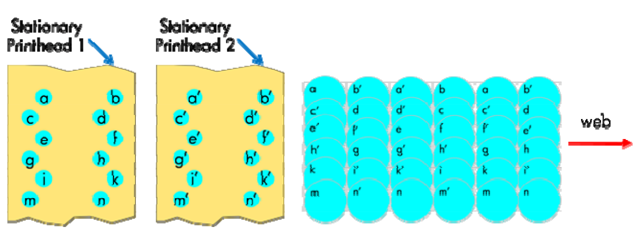

Figure 8 *Nozzle Redundancy example printing 2 drops per pixel*

Aligning nozzles "a'" and "b'" on Printhead 2 with "a" and "b" on Printhead 1 gives four (4) redundant nozzles (e.g., "a", "b", "a'", and "b'") that can print a dot in any given 1200-inch dot-row.

In Figure 8, the neighbors to any dot are seen coming from different nozzle columns and different printheads. This arrangement of dots is continuously randomized during printing to further suppress any periodic patterns resulting from using a particular combination of nozzles. This process is called nozzle cycling, and it provides suppression of nozzle errors comparable to multiple-pass printing in a single pass of the paper under the printheads.

Each Color Plane Processor PCA processes 275 million halftone pixels per second and generates 2 Giga bits of real-time nozzle firing data per second that is delivered to four of the printheads through a fiber optic link.

The T300 web press requires 20 Color Plane Processor PCAs per print side. The PCAs are arranged in an industry standard compact PCI rack.

Printhead driver

At the receiving end of the optical fiber resides the Printhead Driver PCA that is located inside the printbar. This PCA de-multiplexes the incoming data stream and interfaces to each individual printhead die.

Printhead power and control signals are also provided.

Firmware

The modular hardware just described is driven by a scalable firmware architecture that is also modular in nature. Several different C++ components control the different steps of the data pipeline and manage the printheads themselves. These software components are then replicated to match the actual hardware configuration.

Summary

The combination of modular printbars and imaging hardware and software components enables building digital presses with different writing system configurations that can serve different applications with minimal design effort. All it takes is to replicate the basic building blocks.

An example is shown in Figure 9 for a T300 Inkjet Web Press in a 30 inch full color and duplex configuration.

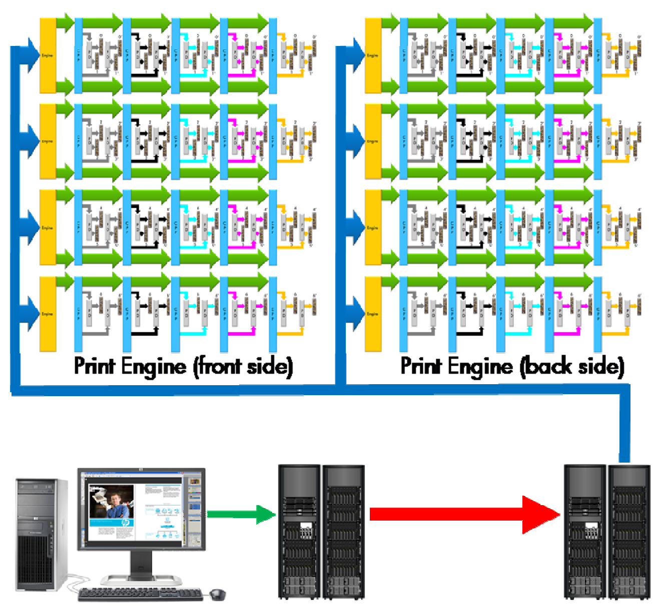

Figure 9 *Components used by the T300 in a full Color Duplex configuration*

References

[1] L. Abello, “Image processing in printing systems” - US Patent Application 10/513,007 Publication number: US 2006/0120787 3.
[2] L. Abello, “Print Engines – US Patent Application 10/513,008 Publication number: US 2005/0260021 A1

Author Biography

Lluis Abello received his MS in computer science from the Polytechnic University of Catalunya (1989). He joined Microsoft as a hardware engineer where he worked on WindowsNT and other research projects. Lluis joined the Large Format printer division of Hewlett-Packard in 1995 and has since then developed hardware and software imaging systems for commercial inkjet products. Currently he is the data path architect of the Inkjet High speed Production Solutions team that developed the Inkjet Web Press family of products.