



HFTA-08.0: Receivers and Transmitters in DWDM Systems

The rapidly growing internet traffic demands a near-continuous expansion of data-transmission capacity. To avoid traffic jams on the data highways, network providers need a technique that provides fast, flexible, and cost-effective bandwidth expansions. One such technique is the data-transmission technology called Dense Wavelength Division Multiplexing (DWDM), which augments network data throughput via the existing fiber infrastructure.

DWDM Technology

In conventional long-haul fiber-transmission systems, data is transmitted at a certain bit rate, using (for low signal dispersion) a single wavelength from the second optical window (1300nm range), or (for low signal attenuation) a single wavelength from the third optical window (1500/1600nm range). To achieve higher transmission capacity, you can raise the bit rate based on Time-Division Multiplexing (TDM), or install additional fiber cables in parallel with the existing ones, or both.

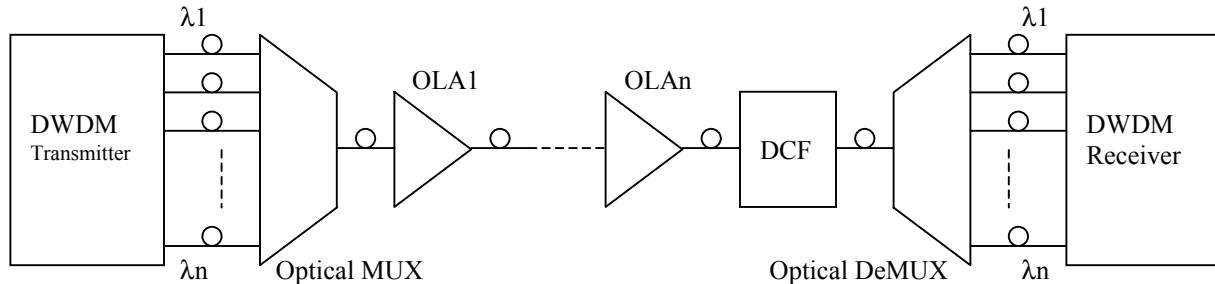
The second approach requires expensive and time-consuming construction work, therefore increasing the bit rate would seem to be the more cost-effective way to achieve higher bandwidth within an existing fiber network. The absence of technologies like a mature and cost-effective process for high-speed IC development, and the limitations of physical fiber media (like fiber-polarization mode dispersion) do not allow the realization of practical commercial transmission systems beyond 40Gbps. Upgrading a single fiber link from 2.5Gbps (for example) to 10Gbps quadruples the bandwidth capacity, but a transmission technology called Dense Wavelength Division Multiplexing (DWDM) can multiply the capacity by multiples as high as 160.

DWDM takes advantage of a physical phenomenon that allows multiple wavelengths of light to travel simultaneously through a single-fiber cable. It allows multiple high-bit-rate signals to ride through the fiber media together, each on a different color of light.

Another significant advantage of WDM transmission (compared with TDM long-haul trunks) is “bit-rate transparency” as conferred by the purely optical functions that are mandatory in such systems— like optical multiplexers and demultiplexers, optical line amplifiers (OLAs), and in future, optical 3R regenerators for ultra-long-link distances. In principle, therefore, the link includes no bit-rate-limiting elements that would require a change of optical line components to achieve a higher bit rate.

Overview of DWDM system components

The basic elements of a DWDM transmission system are the optical multiplexer, the optical line amplifier (OLA), and the optical demultiplexer (Figure 1).



OLA : Optical Line Amplifier (EDFA)
DCF : Dispersion Compensating Fiber

Figure 1. Example of a Dense Wavelength-Division Multiplex (DWDM) link.

An optical multiplexer combines all the received wavelengths of the L band (1530nm-1565nm) and C band (1570nm-1620nm) into one wavelength-multiplexed light signal. Today's systems achieve wavelength separations of 0.4nm or less, which allows about 160 potentially usable wavelengths. The L- and C-band limitation is determined by the optical line amplifier, which is able to amplify incoming light signals from the L- and C-bands only. Optical line amplifiers for the 1300nm window are still under development.

One of the most widely used technologies for implementing the optical line amplifier is the erbium-doped fiber amplifier (EDFA). An EDFA contains a pump laser working at 980nm or 1480nm, which raises electrons to a higher energy level. If light is received with wavelengths within the L or C band, those electrons fall into a lower energy band after emitting photons with the wavelength of the incoming light. The resulting light-domain amplification is independent of the bit rate. Depending on the distance between optical mux and demux, several EDFAs can be cascaded with a typical span of about 100km. This technique allows optical transmission links of several hundred kilometers without need for an electronic signal regeneration.

A disadvantage of EDFAs is that some electrons at the higher energy level generate uncorrelated optical noise by falling into the lower energy band spontaneously. Because DWDM links usually contain a chain of OLAs, this optical noise is amplified in the following EDFAs, and the resulting noise accumulation lowers the receiver's signal to noise ratio (SNR) as compared to systems without OLAs. What's more, this optical noise is asymmetric because it affects logic high levels more than the low levels.

At the receiver side, an optical demultiplexer converts the incoming wavelength-multiplexed signal into the corresponding individual wavelengths launched at the transmitter side. This

demultiplexing function includes very narrow optical filters, for which smaller wavelength separations require greater design effort. Apart from the basic system elements mentioned above, a DWDM system may include other functions—like an optical booster after the optical mux, dispersion compensation, or an optical preamplifier in front of the optical demux—with the aim of improving system performance and extending the link length.

A bit-rate-transparent network (an all-optical network) requires, in addition to transparent DWDM point-to-point connections, additional network elements like optical add and drop multiplexers (OADMs) and optical cross connects (OXC). Available prototypes can demonstrate the feasibility of this purely optical functionality, but today's network equipment (even those called OADM and OXC) contain mainly electronic rather than optical core functions.

Further, ultra-long-haul point-to-point connections may need (depending on line distance) an electronic 3R regeneration in the absence of mature, purely optical replacements. The all-optical network, therefore, is still several years away. Regardless of whether complete or partial all-optical networks are available, though, a network's line terminations must still convert light into electrical signals, because equipment beyond the optical world still relies on electron-based communications.

The network terminations for a DWDM long-haul point-to-point transmission system can be realized with a dedicated line-termination card or with a wavelength transponder. Line-termination cards are used for new installations, where (for example) a central office (CO) transmits to and receives directly from the DWDM link. A wavelength transponder, on the other hand, is essential if a DWDM link must be connected to existing CO equipment that includes the old “noncolored” optical-network interface. The following discussion, valid for line termination cards as well as wavelength transponders, focuses on specific design challenges associated with an O/E receiver and a transmitter in the DWDM fiber network.

DWDM Transmitter

Two features are important for a DWDM system. First, to reduce system cost, the link should be as long as possible without the need for electronic signal regeneration. Second, the system should provide highly reliable data transfers. To improve service quality and extend line distance, a forward error correction (FEC) function can be introduced (see Figure 2).

For pure SDH/SONET data, spare bytes in the signal's frame structure can implement the “in-band” forward error-correction function. Bytes required for the FEC function are inserted into the frame by the overhead-processing ASIC. For protocol-independent DWDM systems, an “out-of-band” FEC must be applied, which increases the bit rate but also increases the efficiency with respect to an in-band FEC. The Reed Solomon FEC algorithm defined in the ITU-T G.975 recommendation is one example of a possible out-of-band FEC implementation. To provide the overhead necessary for the correction function, that algorithm increases the transmission bit rate by 7%.

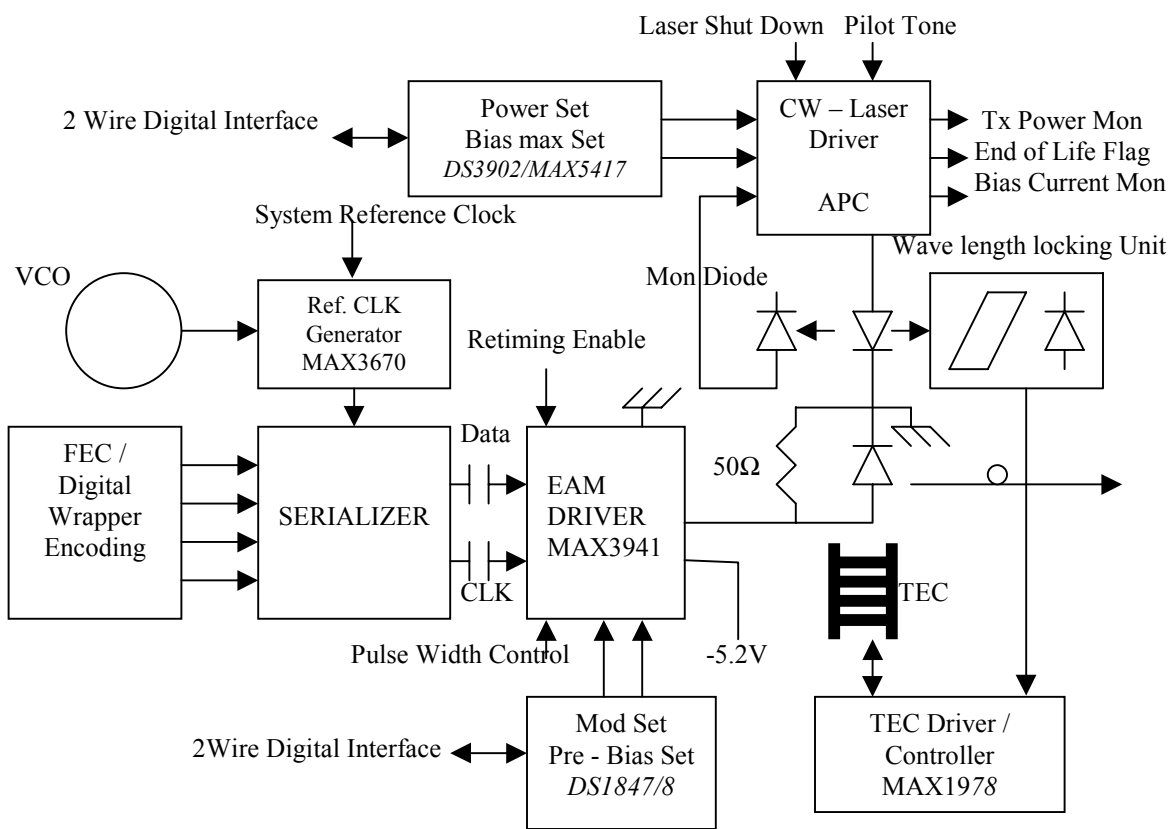


Figure 2. Example of a 10Gbps DWDM transmitter.

Instead of a Reed Solomon FEC, the digital wrapper function defined in ITU-T G.709 is likely to become the champion. The signal is wrapped regardless of bit rate and protocol by a “super frame” that includes (in addition to bytes for the FEC function) the addressing bytes necessary for signal routing—transmitting the payload to its destination. The digital wrapper function’s overhead increases the transmission bit rate by a certain percentage, which in turn depends on which digital-wrapper concept chosen. Regardless of the selected out-of-band FEC/digital wrapper methodology, an additional IC is needed to support the related algorithm, or that function must be integrated into the transmitter’s overhead-processing ASIC.

The FEC or digital wrapper processing is performed on the transmission signal’s lower-speed parallel data stream. Parallel data leaving this processing function, therefore, must be serialized to form the high-speed transmission signal. That task requires a serializer with on-chip clock synthesizer for generating the transmission clock.

For long distance trunks it is very important to launch a low-jitter signal, meaning that jitter generated by the serializer should be as low as possible, as should jitter of the external reference clock applied to the integrated clock synthesizer. In many cases the available system-reference clock not only doesn’t fulfill these jitter requirements, its frequency is also lower than that required. Clock generators with external VCXO or VCSO are available to provide the necessary

low jitter-reference frequency, and fully integrated circuits with internal VCOs are being developed to reduce space and cost.

Because the serializer's output stage is not able to drive an optical transmitter, a driver function is needed. That function adds jitter, unfortunately, so a retiming flip-flop should be integrated into the driver's input stage to minimize data jitter. Usually the serial clock from the serializer is applied to this retiming function, but a non-ideal interconnect between serializer output and driver-retiming input can degrade the clock signal, which can also degrade the transmit signal's jitter performance. The retiming function should therefore be optional.

Another function useful for integration with the driver is pulse-width correction, which introduces a pre-distortion for compensating the non-symmetrical rise and fall transitions in an optical component.

Finally, the serial signal must be converted to an optical signal of dedicated wavelength. For handling up to 160 different wavelengths, the wavelength separation must be no greater than 0.4nm. That calls for an optical source with highly accurate wavelength-stability control, very narrow spectral line widths, and low chirp (the phenomenon of spectral line hopping due to high-speed modulation). Instead of direct-modulated laser diodes, electro-absorption modulators (EAM) or Mach Zehnder modulators (MZ) in combination with CW lasers fulfill the above requirements for long distance transmission.

Housed in modules, these transmitters contain a Peltier element for adjusting to specific wavelengths by setting the temperature, a laser diode that emits continuous light (CW laser diode, DFB type), and a high-speed voltage-driven modulator. The Peltier element (a thermo-electric cooler, or TEC) requires a driver circuit able to handle several amperes for setting the CW laser diode to a specific temperature-related wavelength. To keep an adjusted wavelength constant, temperature must be precisely controlled by the TEC controller circuit.

The TEC controller circuit can be space consuming if all functions must be realized with discrete components like power FETs and operational amplifiers. Fortunately, space-saving and fully integrated TEC drivers with on-chip power FETs and control loops are available to support space-sensitive module integrations and applications with a multi-channel network interface. In addition, a wavelength-locking function is needed for DWDM systems whose wavelength separation is 0.4nm or less, and (depending on the system setup) for 0.8nm separation as well. An etalon-based control unit (Fabry-Perot filter) can keep the wavelength within the tolerance window, with the help of the TEC driver/controller function.

Another important transmitter parameter is the initial user-defined optical-transmit power, which the CW laser must maintain despite aging and variations of temperature. The slope of a CW laser's characteristic curve degrades with time and increasing temperature, so the laser's driver circuit must set and maintain an average optical transmit power. That power level can be ensured by an automatic power-control loop that compares the received photocurrent detected by the CW laser's monitor diode (proportional to optical output power) with an initial defined reference value corresponding to the desired optical output power. In addition, the driver should include an

alarm flag indicating the laser's end of life, a shutdown function for laser safety, a monitor output for the CW laser's bias current, a limit setting for the maximum laser bias current, and an optical average power monitor. Further, a low-speed pilot tone is useful for amplitude modulating the optical output signal. That feature enables (for example) channel identification in DWDM systems.

A modulator driver rather than a direct modulated laser driver should be used to drive an EAM or MZ device, because optical modulators (unlike laser diodes) are usually matched to an impedance of 50Ω . The modulator driver should therefore be optimized for 50Ω loads, and should deliver a modulation voltage rather than a current. EAM devices require a maximum modulation voltage of $\sim 3V$, and MZ types need up to $7V$. MZ modulators provide the narrowest spectral line widths, but require a relatively high modulation voltage, and are more expensive than EAM types. MZ modulators are therefore used in applications that involve ultra-long-haul distances.

Both types require a dc pre-bias of the modulation voltage to optimize the optical modulator's chirp effect. Modulator drivers with internal pre-bias require just one interconnect between the driver output and modulator. That feature allows space-saving module integration, and reduces production effort by eliminating the external inductor usually required for setting up a bias-T network.

DWDM Receiver

Because the optical signal for a DWDM receiver is perturbed by nonsymmetrical optical noise (as explained above) in addition to the fiber attenuation and dispersion that affects conventional TDM receivers, the DWDM receiver carries a greater burden. To increase the receiver's input sensitivity, its first element is usually an avalanche photodiode (APD), which multiplies electrons via a voltage-controlled avalanche breakdown during the conversion of photons into electrons. In order to achieve the multiplying effect, the APD must be reverse-biased (depending on type) up to $90V$.

The reverse bias for an APD must be tightly controlled to keep its multiplication factor (the gain factor "M") constant over temperature. This requires a low noise, low ripple, and highly accurate voltage supply, which should derive the APD's high reverse-bias voltage from the board's available supply voltage ($3.3V$ or $5V$).

To maintain constant gain in an APD, it can be temperature-controlled with a Peltier element, or its reverse bias can be changed as a function of temperature. The second approach is usually more cost effective. An available low-noise bias supply for APDs (an IC) is highly accurate, produces voltages up to $90V$, and includes features such as current limiting for APD protection, an avalanche indicator flag, and an optional DAC for setting the reverse bias.

System management requires detection of the received signal's average power. That can be accomplished right after the APD, in the first pre-amplifier stage (the transimpedance amplifier, or TIA), but the TIA's part-to-part tolerance excludes this approach as the most accurate way to measure receive power. A better way is to detect the average photocurrent directly, from the

photo detector's bias-voltage source. A small current-monitor IC is available for PIN diodes and APDs, which provides a current or voltage output proportional to the average photocurrent. That product allows accurate detection even for photocurrents below 1uA.

After devising circuitry for the receiver diode, the designer must deal with optical noise launched by the OLAs. Being asymmetric, that noise has a higher noise floor on logic 1 than on logic 0, which reduces the BER significantly in a traditional receiver. As a result, the receiver chain's clock and data recovery (CDR) decision circuit (which distinguishes between logic 1 and 0 by performing a time-and-amplitude decision on the incoming signal) must have the capability to adjust the threshold level of its decision voltage before the amplitude decision is made. That threshold adjustment shifts the amplitude-decision level from the middle of the signal-eye opening towards logic 0, thereby achieving a symmetrical eye opening relative to the decision level.

For a successful implementation of this BER optimization, the incoming signal should not be distorted by electronic functions in front of the CDR. It is therefore essential that the signal-to-noise ratio undergo minimum change between the APD and the decision function. As a consequence, the preamplifier that converts APD current to voltage must implement linear signal amplification over the entire dynamic range, and the following post-amplifier must add further linear amplification without clipping. To facilitate adjustment of the voltage decision threshold, the task can be accomplished by a linear, automatic gain control circuit (AGC) that provides a constant voltage swing at the CDR input over the receiver's full dynamic range.

The adjustment can be performed manually, deriving the decision threshold level from experience, or via an automatic control loop that measures the BER. The manual adjustment is cost effective for low bit rates (to 2.7Gbps), but for bit rates of 10Gbps and higher an automatic BER optimization should be considered, as a consequence of the bit-rate-dependent lower margin of the signal eye opening. If an FEC or digital wrapper decoding function is implemented on the receiver board right after the CDR and deserializer, the actual receiver BER can be derived from this function, which counts the corrected errors on the received signal. Such error-counter information can then be used as a criterion for the feedback loop that controls the automatic threshold-level adjustment (Figure 3a).

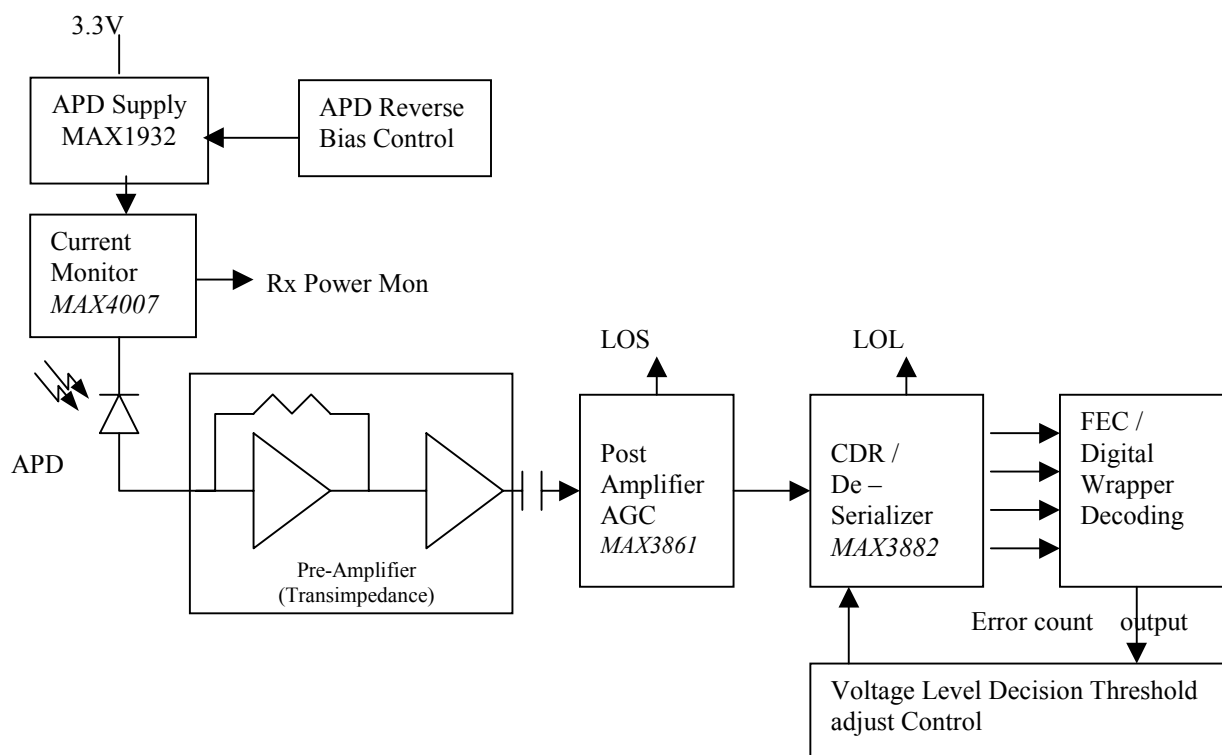


Figure 3a. Example of a 2.5Gbps DWDM receiver with linear preamplifier and AGC

An alternative method of adjusting the threshold level is to control the dc voltage at the preamplifier output. As in the previous method, this entails linear amplification over the dynamic input range of the preamplifier, plus an adaptive, automatic control of threshold level. Because the output amplitude of the preamplifier is not constant, there is no alternative to automatic threshold-level control, which receives its feedback from the FEC or digital-wrapper error-counter output.

An advantage of controlling the threshold at the preamplifier output is that a simple limiting amplifier can be used instead of an AGC function. An amplitude-decision circuit after the preamplifier, like a limiting amplifier, is acceptable because the threshold level for the amplitude decision is defined at the preamplifier output (**Figure 3b**).

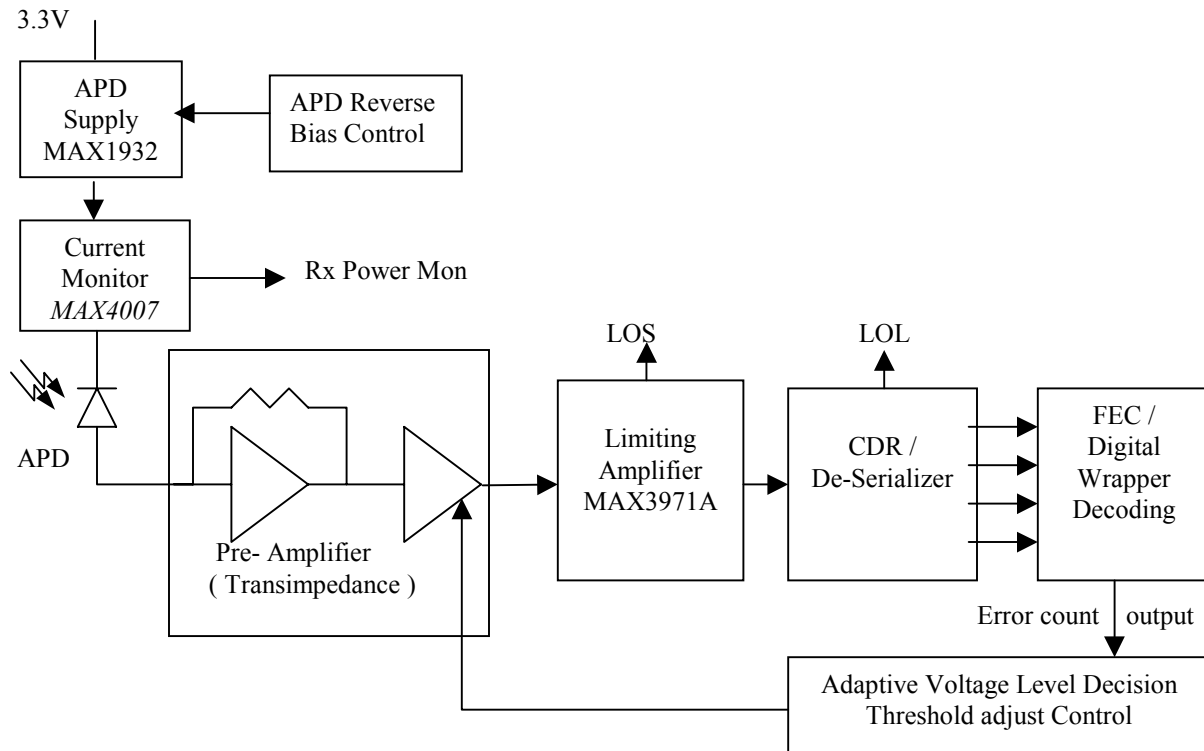


Figure 3b. Example of a 10Gbps DWDM receiver with linear preamplifier and limiting amplifier.

This article first appeared in Communication Systems Design, June, 2003.