Scalable passively mode-locked semiconductor lasers for microprocessor clocking

James P. Siepmann and Adam Rybaltowski
Research & Development Department, LightTime, Oshkosh, WI 54901

ABSTRACT

The integration of photonic clocking in microprocessors is anticipated to occur during the 2008-2012 high-volume manufacturing (HVM) cycle. Though photonic clocking can be achieved through electronic modulation or actively mode-locking a laser, a more cost-effective and better solution would be to use internal cavity passively mode-locked semiconductor lasers. Not only do these lasers offer low-cost, simplicity, and ease of integration, but prototypes that are amenable to HVM are currently available. We present such a laser that is scalable by design to clock rates of 9 to hundreds of GHz and wavelengths in the 800 to 1100+ nm range. These lasers utilize internal saturable absorber(s) to passively mode-lock a semiconductor laser with relatively high peak powers. Experimental results from these lasers show an RF spectrum signal peak that is at least 40 dB above the noise floor with a -10 dB width of <1 MHz. The RMS jitter as determined by an oscilloscope with a precision timebase module was found to be ~1 ps which is among the best for this type of laser. Autocorrelation was used to confirm mode-locking and pulse width. In addition to experimental data, a theory and discussion on how the different characteristics of these lasers can be tailored for various commercial applications such as microprocessor clocking will be presented.

Keywords: mode-locked semiconductor laser, diode laser, passive mode-locking, microprocessor clock, photonic clocking, saturable absorber, optical pulse multiplier, laser theory, ps pulse generation.

1. INTRODUCTION

Photonic clocking and signal distribution in multi-gigahertz microprocessors hold the promise of decreasing jitter, skew, delay, and power consumption while traveling longer distances and maintaining the signal's integrity. Much like fiber optic technology revolutionized the telecom industry, a promising high-speed/low-cost laser technology may usher in a new wave of possibilities. The photonic clocking of microprocessors is inevitable; the only question is when it will occur.

Problems with keeping the current copper based electrical interconnect are anticipated to begin in 2007[1] and Intel anticipates its switch from electrical to optical to occur during a 2008-2012 high-volume manufacturing (HVM) cycle.[2] This switching from electrical to optical will entail the implementation of on-chip optical waveguide interconnects and an optical clock to provide the signal source. With ITRS (International Technology Roadmap for Semiconductors) anticipating on-chip clock rates in excess of 10 GHz by 2008 and over 20 GHz by 2012, a scalable photonic clocking system will be needed.[1]

One of the more disruptive applications for this laser technology may be in redefining how microprocessor clocking and signal distribution (in both single core and the emerging multi-core/high traffic environments) transition from voltage controlled oscillator (VCO) quartz-based clocking to optical clocking, at native, multi-gigahertz speeds. While microprocessors can and will be driven to higher performance with the incumbent technology, the "clock is ticking" when it comes to prolonging the tenure of quartz based technology and electrical signal distribution. While the transition to optical clocking is generally thought to be "when, not if", the computer microprocessor market segment being highly price sensitive will only welcome a photonic solution if it can improve performance while keeping costs relatively low.
Though temporary or transitional technologies such as optical pulse multipliers may allow the use of inexpensive modulated lasers to achieve high frequency outputs (e.g. 2 GHz to 20 GHz)[3], it is still less expensive in HVM to use an internal cavity passively mode-locked laser. The use of optical pulse multipliers though may serve a useful purpose as a distribution system in and between microprocessor chips.[3]

Physics and Simulation of Optoelectronic Devices XIV, edited by Marek Osinski, Fritz Henneberger,
Yasuhiko Arakawa, Proc. of SPIE Vol. 6115, 611504, (2006) · 0277-786X/06/$15 · doi: 10.1117/12.640882

As the clock rate for microprocessors increases, certain problems become more pronounced with electronic clocking. For example, in traditional microprocessor transmission lines at multi-GHz frequencies, there is a tenfold decrease in a signal's strength with every 2 gigahertz (GHz) increase in electrical signal bandwidth.[4] Not only will it require significantly more power to increase the clock rate and its distribution, as microprocessor copper interconnects are further scaled down (as outlined in the International Technology Roadmap for Semiconductors (ITRS) guidelines[1]), the distance that they can transmit a signal becomes very limited and fraught with increasing delays. Figure 1 shows the delays associated with electrical, optical using a transimpedance amplifier (TIA), and receiverless (photodetector only) optical systems[5] while Figure 2 puts them in perspective relative to the clock cycle. Both figures are based upon using a waveguide for the optical distribution but a focused free-space distribution of an optical signal[6] would have even less delay. It becomes readily apparent that there will need to be a switch to photonic clocking in the near future, especially when electrical clock skew and jitter threaten to consume the majority of the clock cycle as shown in Figure 3.[6] As Intel's platform group noted, "At higher frequencies, clock skew and jitter and latch delay become a much bigger percentage of the clock cycle, reducing the percentage of the clock cycle usable by actual logic."[7]

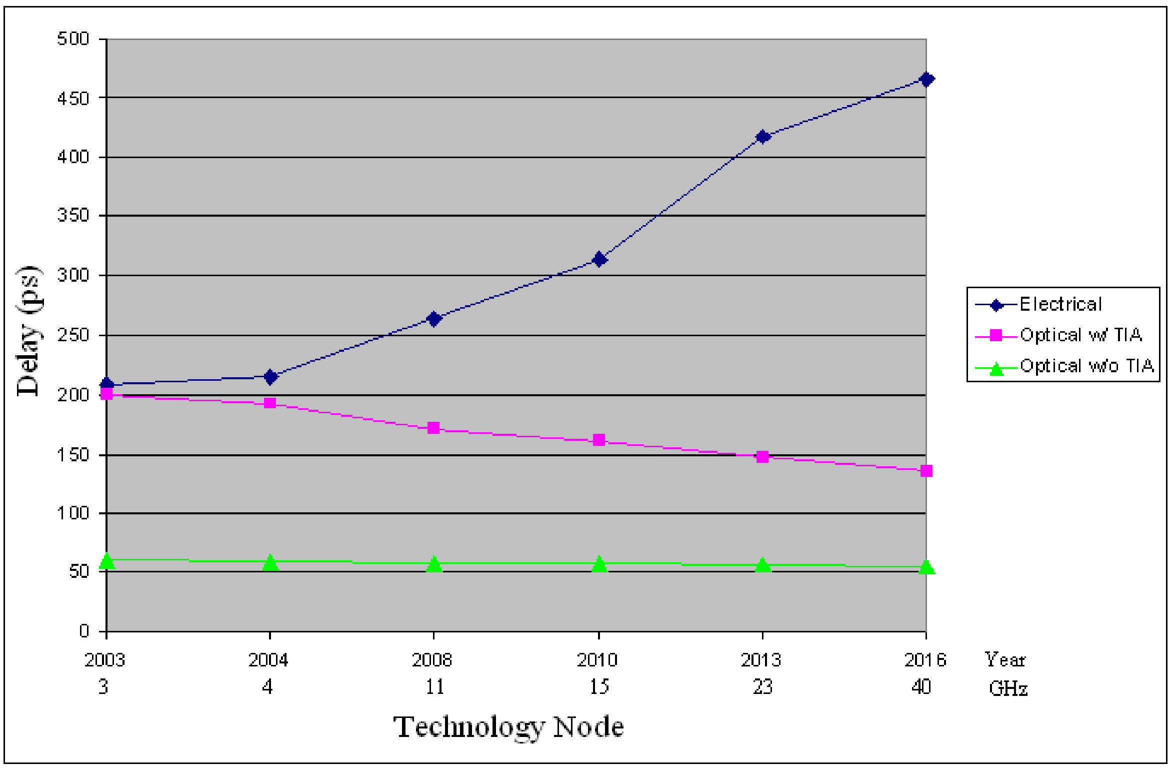

Figure 1. Delay associated with a clock signal traveling 1cm in a microprocessor. The electrical distribution is calculated for global interconnects (wires) of scaled copper per ITRS guidelines[1] and the optical signal distribution is calculated for a waveguide.

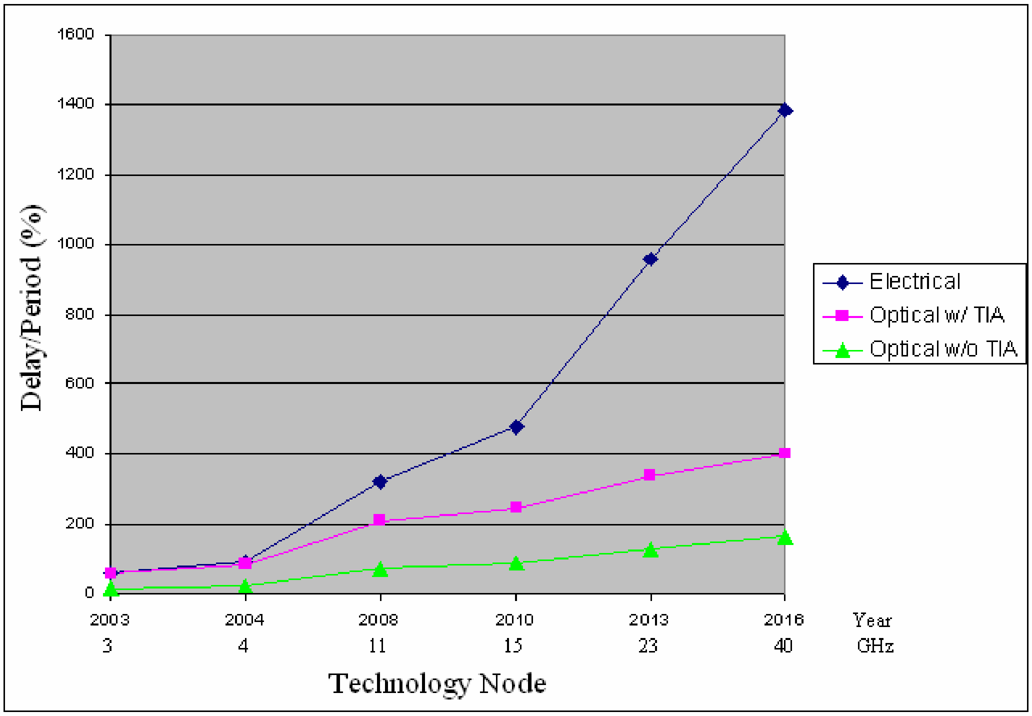

Figure 2. Delay associated with a clock signal traveling 1cm in a microprocessor as a percentage of the clock cycle. The electrical distribution is calculated for global interconnects of scaled copper per ITRS guidelines[1] and the optical signal distribution is calculated for a waveguide.

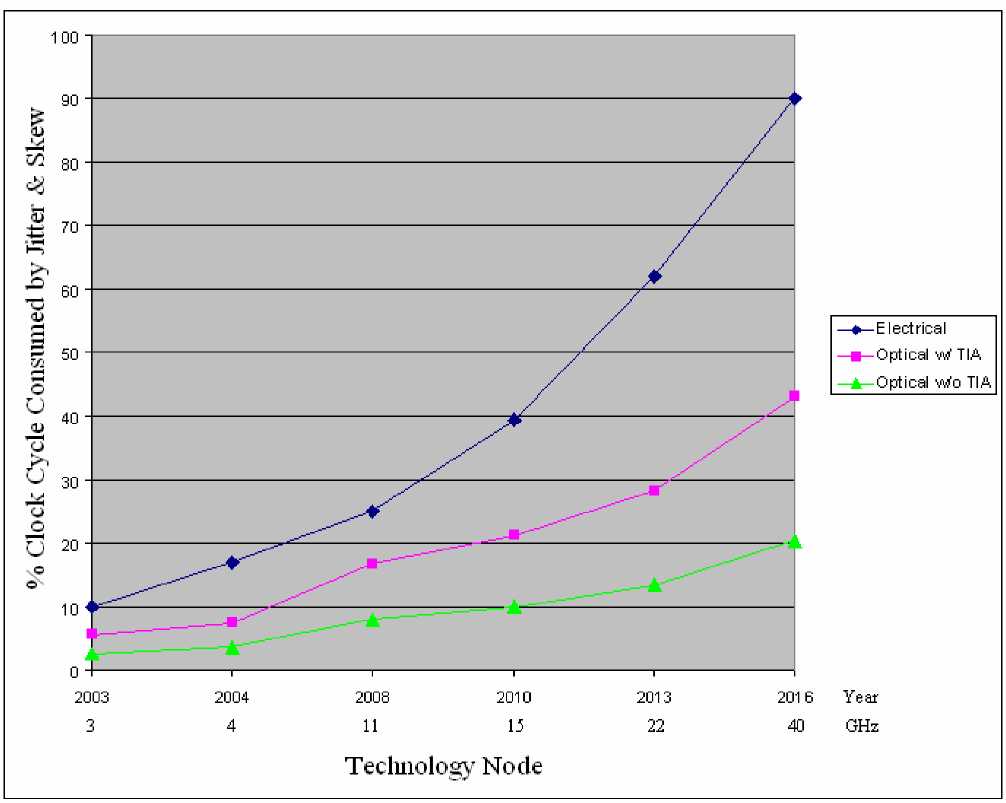

Figure 3. Actual and extrapolated combined jitter and skew as a percentage of the microprocessor's clock cycle.[6]

2. POSSIBLE MICROPROCESSOR CLOCKING SOLUTIONS

One option for extending the useful life of copper in signal distribution is to use repeaters, but this will consume more power and chip space as well as make the interconnect routing more difficult. Another option is to not scale down the clock copper lines, but that will only decrease the bandwidth, significantly influence on-chip transients, increase crosstalk by up to 60%, and cause other problems.[8] Additionally, with the move to multiple core processors, the increased distances involved in inter-chip as well as intra-chip transmission may soon be too great for copper. These challenges along with electromagnetic interference, voltage isolation, impedance mismatch, inductance problems, electromigration, edge integrity degradation, and more that are inherent in electrical clocking, increasingly point to photonics as the only likely choice for the next generation of microprocessors.

Photonic clocking, simply put, consists of a laser or light emitting diode (LED) that generates a high frequency optical signal that can then be distributed to distant sites by an optical waveguide or fiber. Currently, this signal will need to be received by photodetectors, which then generate an electronic signal for the microprocessor.

The generation of an optical clocking signal can be accomplished by direct gain modulation or by turning the output beam of a continuous-wave laser or LED on and off using an optical modulator. Both methods of modulation though are limited in frequency/have relatively slow signal rise and fall times, and require additional electronics which increase power consumption, heat generation, and cost.

There has been much press this year about “silicon lasers” which basically are a Raman amplifier that still requires an optical pump such as a semiconductor laser thereby defeating the monolithic integration advantage it may have had over using a passively mode-locked semiconductor laser. There are attempts to electrically pump Er-doped silicon nanoclusters but reported gains have been very low and there are many other challenges associated with this method.[9] Even if silicon lasers should become practical, they will require modulation (either of the laser or its output), which will increase their cost and complexity relative to passively mode-locked semiconductor lasers.

Unlike an LED, a laser can produce a pulsed beam by “mode-locking” its multiple cavity modes to yield repetitive pulses of high peak power and short duration. Many of these mode-locked lasers are too large and costly though to be practical (i.e., solid-state lasers, fiber lasers), semiconductor mode-locked lasers however are a viable option.

Semiconductor lasers that mode-lock come in two basic forms: actively mode-locked and passively mode-locked. Actively mode-locked semiconductor lasers require an external RF source, which is not practical at high GHz rates and would be cost prohibitive in this context. Passively mode-locked semiconductor (PMLS) lasers, which by definition do not require an external RF source, are a more viable low-cost photonic clock solution for microprocessors.

Of the PMLS lasers there are two sub-types: internal cavity and external cavity. The external cavity versions though potentially quite accurate, require multiple components that need meticulous placement thereby increasing the manufacturing costs dramatically. One company that developed such lasers for the telecom industry was charging around $50,000 per unit. Alternatively, the internal cavity lasers reported in this work are a relatively inexpensive alternative and amenable to HVM. For the rest of this paper, we will be dealing with internal cavity PMLS lasers and will refer to them as simply PMLS lasers with the acknowledgement that they are of the internal cavity type.

3. PMLS LASER DESIGN

3.1 Mode-Locked Operation of PMLS Lasers

Mode-locking is a means of generating very short pulses from lasers that have pulse widths below those typically available by direct modulation of the laser or its output. In mode-locking, instead of shaping the envelope of the laser field directly, longitudinal modes of the laser are coupled, that is, the fields of each longitudinal mode are added (more or less coherently), which results in the modulation of the total field with the beat frequency and a short-pulsed output. The main advantage of this technique of generating short pulses over direct modulation is that coupling of the modes does not require a bandwidth comparable to that of the output pulses; in fact, a DC input may be sufficient. As a result, costly and bulky RF or microwave components can be eliminated. Furthermore, shorter pulses, higher pulse repetition frequency, and higher peak power can be achieved in mode-locked semiconductor lasers.

The details of mode-coupling techniques vary depending on the type of laser and mode of operation. In semiconductor lasers, these techniques can be divided into passive, active, and hybrid mode-locking. In active mode-locking, an external signal is applied at the frequency of natural output pulse repetition or its (sub)harmonic. In passive mode-locking, no modulation input is provided and the laser pulses at its natural repetition frequency. Hybrid mode-locking is a combination of the two, where RF is applied to either a saturable absorber (SA) section or an independent section at the output end of the laser. In a hybrid design, subharmonics could be used thereby decreasing the cost and complexity of the additional electronics. Finally, optical injection could be used instead of RF input to stabilize a hybrid version at the desired frequency.

To implement mode-locking in semiconductor lasers, a SA is typically used. This section is physically the same as the gain section (GS), but it is not forward biased to provide gain. Instead, it is typically reverse-biased, which results in absorption rather than gain and most importantly, this absorption is saturable. The reverse-biased SA sections in a laser cavity provide a bleachable medium for mode-locking. In the time-domain description of mode-locking, two counter-propagating pulses circulating in the laser cavity experience lower absorption when they overlap in the bleachable medium (SA section), thus favoring pulsed operation over continuous-wave operation in the laser. By varying/adjusting the reverse voltage across the SA section, the density of carriers capable of absorbing light is varied, which leads to the change of the saturation parameter controlling the strength with which lower-power light (pulse tails) is discriminated against higher-power light (pulse peak) by stronger absorption. This different absorption of weaker portions of the time-varying beam leads to a light power instability, with a positive-feedback mechanism favoring higher transmission of (already) higher-intensity light. In other words, a light peak is absorbed less than its tails, and this is sufficient to shorten existing pulses, or to spontaneously generate ps pulses of light in a properly designed semiconductor laser.

Typically, there is an optimal SA length in internal cavity PMLS lasers. When an SA is too short, it does not perform its function of absorbing light with a sufficient strength to narrow the light pulses. This is because absorption requires a certain product of the absorption coefficient (which is independent of the length of the absorbing material) and the absorption section length. Since the former is practically fixed by the laser/GS design and external bias, both of which can not be varied much, the most efficient (if not the only) control of absorption in this class of lasers is by varying the SA relative length. On the other hand, the SA length can be too large, for the following reasons. First, with the increasing SA absorption, not only the tails of each pulse but also - to some degree - its peak is absorbed, which decreases the output power. Second, with a fixed pulse repetition rate and thus the total laser length, the increase in SA length can only be realized at the expense of the GS length, which leads to a decreased output power and efficiency of the laser. Finally, if the length of the SA is increased too much then the device will self-pulsate and not mode-lock.

3.2 Effect of Output Facet(s) Reflectivity

For the purposes of this paper AR will be a coating that results in 2-30% reflectivity (albeit typically it is only a few percent), HR will be a coating that results in 70-98% reflectivity (of which an uncoated laser facet would fall into this category since it has about a 70% reflectivity), and a mirror coating will have ~100% reflectivity. Though coatings can be used to yield a reflectivity in the 30-70% range, the distinction begins to blur the closer the reflectivity gets to 50% in terms of SA placement and design for optimal mode-locking. Then for any given laser the effective cavity frequency will be $c/(2d\eta_g)$, where c is the speed of light in vacuum, d is the physical cavity length, and η_g is the group refractive index, while for a version with a single mirrored facet (~100% internal reflectivity) it will be $c/(4d\eta_g)$. This is assuming that the output facets of the laser are uncoated, while if an AR (anti-reflectivity) coating is used then the formulas are $c/(d\eta_g)$ and $c/(2d\eta_g)$ respectively. This is because the HR coated laser with its high reflectivity effectively doubles the cavity length for the purposes of modelocking while for an AR coated laser this effective cavity length is the same as its physical length. For instance, a 4mm laser that is AR coated at both ends with an η_g=4.15 would have an effective cavity frequency of about 18GHz while an HR coated laser of the same length will have an effective cavity frequency of 9GHz.

In terms of optimal mode-locking, the modes overlapping in the SA should be of sufficient and similar power. In the case of a centrally located SA, that is not a problem and is optimal regardless of the reflectivity coating, but for an SA that is adjacent to the output facet, the situation is quite different. In the latter situation, the pulse entering the SA on its way to the output facet will be of maximum power while the returning pulse, which is effectively the overlapping mode from the opposite direction, will be of a diminished power. This situation is critical for an AR coated laser while not being significant for an HR coated one, which can be illustrated by using a simplified version of the cavity power dynamics where reflectivity is the only factor. For instance, a laser that has a facet with 80% reflectivity will have roughly four times the internal cavity power of one with 5% reflectivity and only 20% of its power will be lost at that facet in contrast to 95% for the other. Therefore, the modes in the HR coated laser will not only have a higher cavity power than the AR coated laser but they will be similar (only a 20% difference in power of the overlapping modes), while the AR coated laser will have not only a lower cavity power but there will be a 95% discrepancy between the power of the overlapping modes at the SA adjacent to the output facet. This aspect of reflectivity has been one of the key components missing from prior theories and must be taken into account in the design phase and will be discussed further in Section 3.3.

Facet reflectivity also will have an effect on other PMLS laser characteristics such that as facet reflectivity is increased, there will also typically be an increase laser diode heat generation, a decrease in the average power output, and a decrease in the PW (pulse width). This decrease in PW is due to the increased passes that occur through the gain and SA sections where linewidth is improved. Also as the carrier lifetime difference (as determined by the absorption recovery time) between the gain and SA sections becomes more pronounced (e.g., by increasing the reverse bias of the SA as shown in Figure 4), the PW will also decrease. With this range of absorption recovery times, excellent secondary pulses can be suppressed for mode-locked frequencies of 20GHz or greater. Additionally, with an absorption recovery time that is less than 4ps, these PMLS lasers have the potential for generating frequencies in excess of 250 GHz.

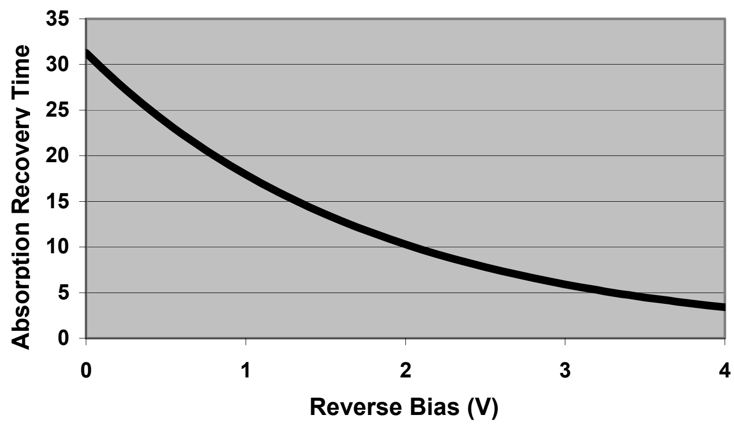

Figure 4. Calculated absorption recovery time plotted against the reverse bias of the SA using an experimentally derived fit for the lasers tested.[10]

3.3 Design and Placement of the Saturable Absorber(s)

Typically, the non-mirrored version has been called a colliding pulse mode-locking (CPM) laser and the latter a self colliding mode-locked (SCPM) laser, but these are for the most part definitions without distinction as can be seen in the two examples shown in Figure 5. In the case of a mirrored version where an SA is not located adjacent to one of the facets or there are multiple SAs, the terminology is not appropriate. Also because the waves do not technically "collide" but rather overlap in the SA, we will only use the term PMLS for the rest of this paper in lieu of either CPM or SCPM. Though the non-mirrored PMLS laser typically has the advantage of a shorter PW, its dual output is not commercially desirable and the required double cavity length in many cases leads to a decreased output power and slope efficiency, i.e., lower increase of the output power with the same increment of the input current above the lasing threshold. Either way the mirrored and non-mirrored functional near-equivalent designs can be easily extrapolated from each other.

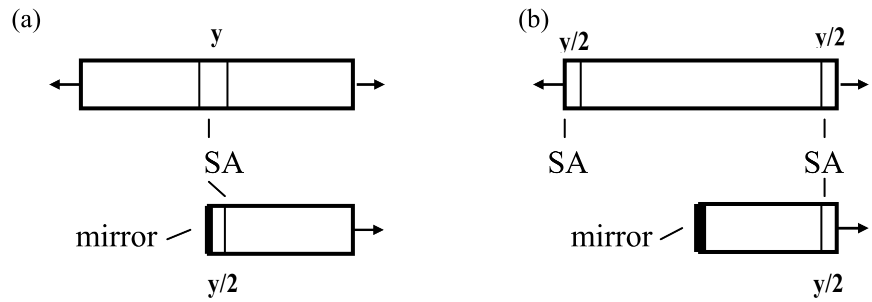

Figure 5. Examples of CPM lasers (above) and their SCPM equivalents (below). The SA length is denoted by the length y or a fraction thereof.

Simply put, mode-locking occurs as the SA(s) subdivides the effective cavity frequency (as earlier defined) into multiple modes which are then locked by the SA(s). In a well-designed laser, this mode will be physically occurring at the center of the SA(s) but in less well-designed lasers this may not be the case. The expectation of how PMLS lasers will work does not necessarily agree with the experimental results as Martins-Filho et al discovered when their 2-SA laser, represented in Figure 6(a), generated three times the effective cavity frequency rather than the expected four.[11] One would have expected the frequency to be four times the effective cavity frequency because the length of the center GS (GSs being measured on center relative to the SAs) or twice that of an end one (since it is effectively twice its physical length due to the facet being uncoated) was ¼ the effective length of the laser cavity as shown in Figure 6(b). From Figure 6(b) it can be seen why the design produced the results it did since there are only three GSs bound by SAs and a gain section only contributes to mode-locking when it is bound by SAs. Physical representation (a) below would not have mode-locked except that the SAs used were relatively large enough to encompass a mode that was one-third the effective cavity length. Figure 6(c) represents a better effective cavity design for a PMLS laser (for both AR coated and HR versions) which would mode-lock at four times its effective cavity frequency. Figure (d) is a physical representation of a HR coated version of the effective cavity length seen in (c).

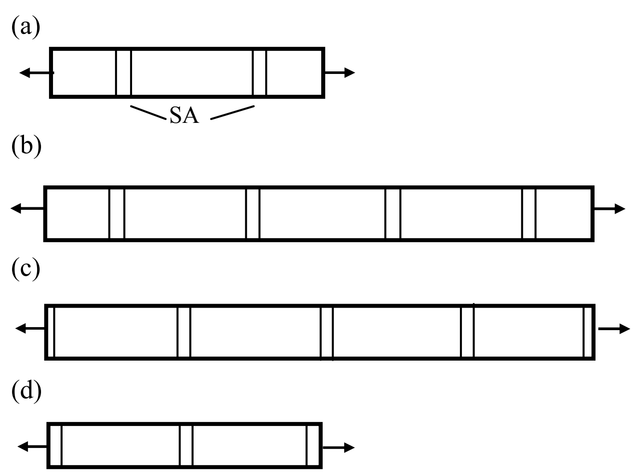

Figure 6. The top-view of a 2-SA non-mirrored PMLS laser is shown in (a) where the physical length of the center GS is twice that of the outer GSs. The effective cavity length with AR coatings would be the same while for a HR coated laser it would be as represented in (b). Figure (c) represents the optimal effective cavity design to yield 4 times the effective cavity frequency for all coatings. While Figure (c) also shows the physical representation for an AR coated laser Figure (d) shows the physical representation for a HR coated laser.

Optimal design rules for a non-mirrored PMLS laser can be summarized as:

1. *For optimal mode-locking of a laser to its effective cavity frequency, a single center SA should be used for an AR coated laser while a single SA adjacent to one of the output facets should be used for a HR coated laser. For all other designs use the following rules.*
2. *A SA should be adjacent to the output facet.*
3. *The length of an SA adjacent to a facet should be ½ the length of those that are not.*
4. *The length of a GS adjacent to a mirror should be ½ the length of those that are not.*
5. *The output frequency of a non-mirrored mode-locked laser will be the effective cavity frequency times the effective number GSs. This corresponds to physical number of GSs in an AR coated laser and twice that in a HR-coated laser.*

Rule #1 allows for an AR coated laser to use what little reflectivity it has to double the effective length of the GS by letting it be bound by a single centrally located SA. Placing the SA at the output end(s) of a laser will not be conducive to optimal mode-locking for the reasons previously stated in Section 3.2 though it can be done, especially with coatings that have a reflectivity greater than the typical AR coatings. It can be seen that these designs though are less desirable than the ones with multiple SAs in terms of mode-locking but there are disadvantages, because if the lasers becomes too long the relative power and pulse quality begin to decrease. It should be noted that for the HR coated version, a near-equivalent mirrored version is not possible because of its asymmetry. Optimal designs for these lasers that are mode-locked to their effective cavity frequency are shown in the first row of Figure 7.

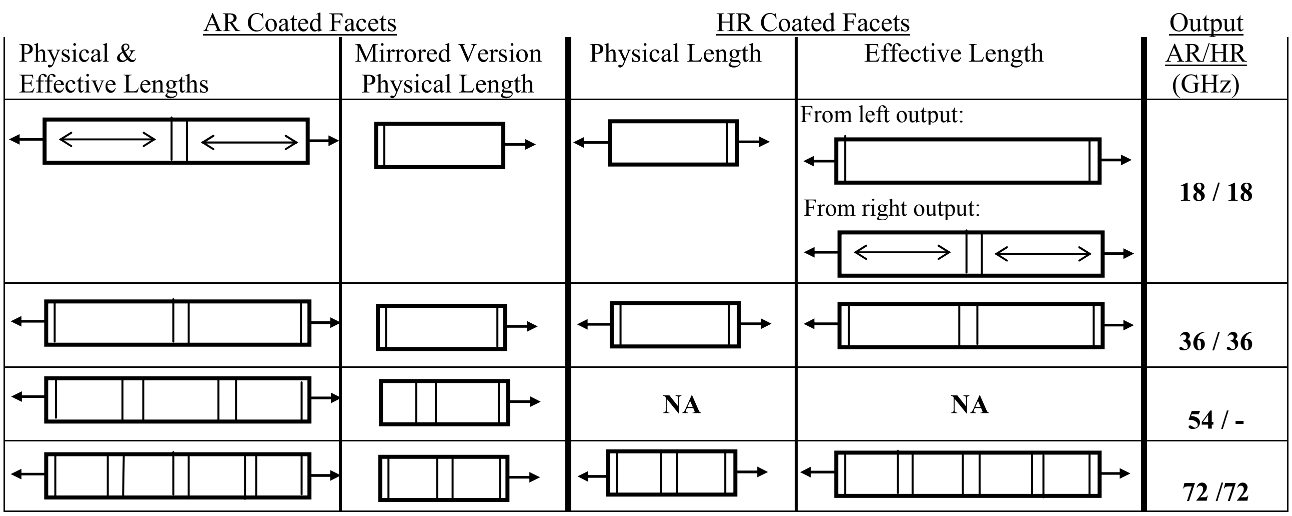

Figure 7. Optimal design examples for non-mirrored and mirrored (~100% reflectivity) AR coated PMLS lasers are shown along with those for HR coated versions. A GS is a forward biased section that is effectively bound by reverse biased SAs. A mirror is located at any facet that does not have a beam (→) coming out of the laser. Note that the effective cavity frequency for all of the lasers is 18GHz along with an effective length of 4mm and η_g=4.15. Their corresponding mode-locked frequency output is noted in the right column.

Using these rules of optimal design, examples of various mode-locked laser configurations are shown in Figure 7. It is interesting to note that there are no optimal designs for HR coated PMLS lasers to achieve odd multiples of the effective cavity frequency. This is because the effective cavity length will always have twice the number of SAs that the physical cavity has and therefore always an even multiple. Though one can graphically convert a mirrored version into its non-mirrored near-equivalent in order to apply these rules, a numerical conversion of the number of GSs in a mirrored version to a non-mirrored can be accomplished by multiplying the number of GSs by two and then adding one more if there is a half GS next to the mirror. This number of GSs can then be multiplied by the effective cavity frequency to determine the actual output frequency.

AR coated PMLS lasers are less effective in mode-locking than their HR coated counterparts since the SA adjacent to the output facet has much less power returning to the second half of the SA as most of the power has passed through the facet to the outside. There are other advantages that may favor their usage though that will be discussed in later sections.

4. FABRICATION OF EXPERIMENTAL PMLS LASERS

Though our lasers were multiple quantum well, the techniques of design can be used for a variety of semiconductor lasers. The active region of the demonstrated lasers consists of four 10nm GaAs quantum wells with 10nm $Al_{0.2}Ga_{0.8}As$ barriers, yielding an emission wavelength of about 860nm. The cladding layers are $Al_{0.43}Ga_{0.57}As$ and doped to 2-3E17 with C (p-type) and Si (n-type). The p-type contact layer is GaAs doped to >2E19. The entire structure is grown on a GaAs substrate.

Since uncoated lasers have been well described and can not be optimally designed for mode-locking at odd multiples of the effective cavity frequency based on our theory, we used an AR coating for the output facet. We also chose to use a mirrored laser with a single output for higher output power and commercial viability. By using a single AR coated facet we were able to achieve higher average power, longer lifetime, and a wider pulse, of which the latter was thought to be more appropriate for the optoelectronic interface of a microprocessor.

We fabricated ridge lasers with several different laser geometries/designs, all being 2mm or 4mm in physical length with a mirrored facet and an AR coated facet. The optimal design for achieving the same mode-locked frequency as the effective cavity frequency (and therefore the lowest possible frequency) was the design that utilized a single SA adjacent to the mirror as illustrated in the top row of Figure 7. Another design that had the SA adjacent to the AR coated facet did not mode-lock but if our facet had been uncoated we would have anticipated finding it optimally mode-lock at the effective cavity frequency as illustrated in the top row of Figure 7. Another design that had the SA centrally located also did not mode-lock for reasons of GS length differences (which our SA was not large enough to compensate) and poor SA placement (not maximally away from the output facet).

Also as the forward bias of the GS is increased to a multiple of the threshold, the length of the SA may need to be decreased in order to avoid self-pulsation. We found that a SA length 25-75μm worked well for most laser designs though larger and smaller lengths would be expected to still work. The gap between the SA and GSs was kept at 10μm because if the gap became too narrow (<8μm), electrical isolation between the sections became a problem. Likewise there was no need to have a larger gap because it would not only reduce the output power of the laser for a specific forward bias but also be acting as a SA with a reverse bias (RB) near 0V.

5. PMLS LASER PERFORMANCE

Our PMLS lasers had low jitter and high pulse energy. These PMSL lasers utilizing our rules can be scalable from 9 to 100+ GHz while its wavelength can range from about 700 to 1600nm by varying the material system used. Lasers have been designed and fabricated at 860nm as well as for 1100+nm. The 1100+nm designs are especially important since silicon is transparent above 1100nm those lasers can be coupled with a properly designed silicon waveguide for an efficiency of over 80% based upon our simulations.

Though our lasers mode-locked over a broad range of settings, optimal functioning was achieved with a forward bias that was 2-3 times that of threshold (typically ~60mA) and a RB in the 0.2V to 1.0V range. This RB is theoretically consistent with what would be expected of a PMLS laser mode-locked at 18GHz based upon the absorption recovery time for that RB range as shown in Figure 4. Though a higher RB still allowed mode-locked operation, it would require the forward bias to be increased in order to compensate for the optical power decrease. Therefore it is always better to use the lowest RB possible.

These PMLS lasers tested were fiber coupled and then optically attenuated prior to introduction into a 50GHz photodetector. The data shown in Figures 8-10 was obtained from a 2mm PMLS laser with a single SA adjacent to a mirrored facet and an AR coated output facet (optimal design per rules of Section 3.3), mode-locked at 18GHz. Figure 9 shows a continuous pulse stream with only 1 ps of RMS jitter using a 50GHz oscilloscope with a Precision Timebase reference module. Depth of modulation was found to be over 90%. Pulse energy was over 3pJ with a peak power greater than 160mW at the designed pulse width of 19ps FWHM. Both low and high-resolution RF power spectrum graphs shown in Figure 8 were obtained using a 40GHz spectrum analyzer and an autocorrelation tracing is shown in Figure 10.

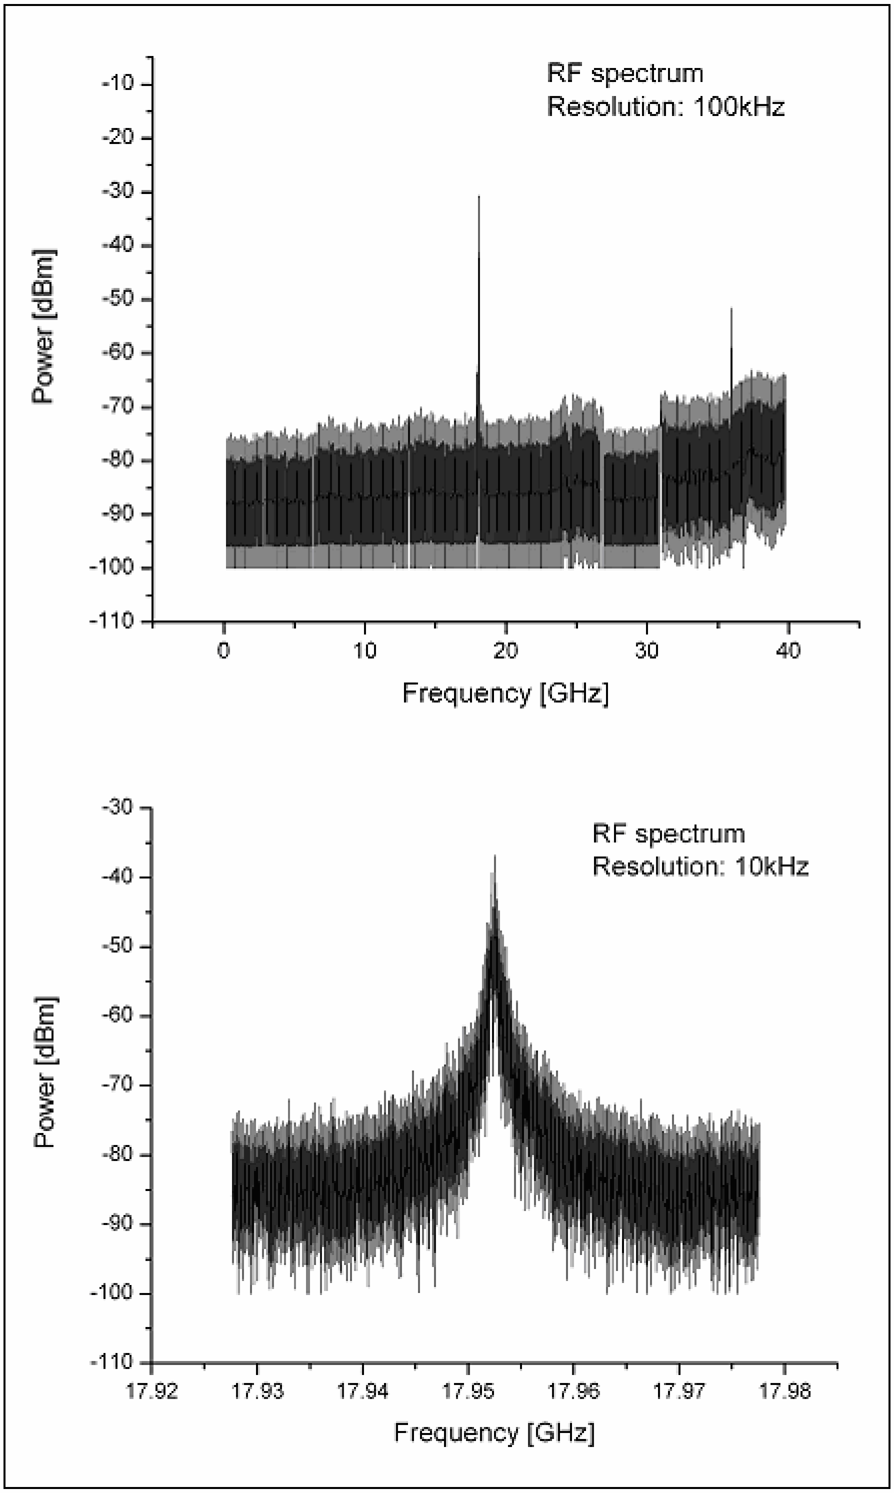

Figure 8. The low-resolution RF spectrum shows no self-pulsations or extraneous pulses except for the demonstrated second harmonic. The high-resolution graph shows a narrow –10dB peak width that is <1 MHz. Data was obtained using a 50GHz photodetector in conjunction with Agilent's 8564EC 40GHz spectrum analyzer.

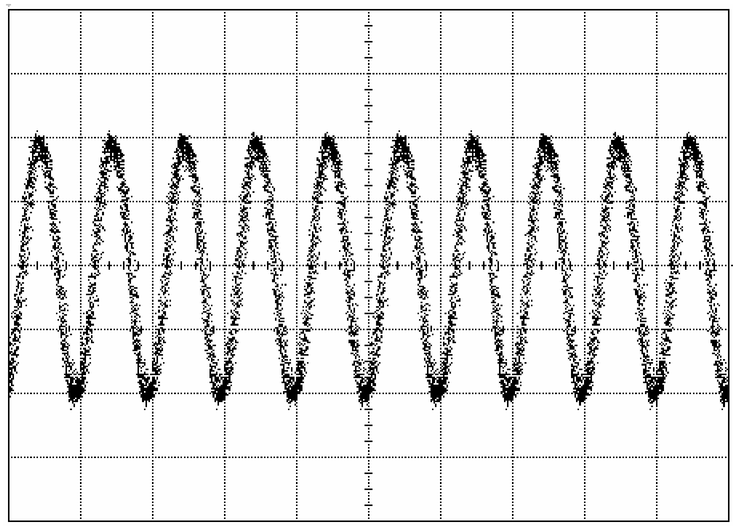

Figure 9. Pulse train from a PMLS laser using a 50GHz photodetector and an Agilent 50GHz oscilloscope with the 86107A Precision Timebase module. Mean RMS jitter was 1.153ps with a standard deviation of 11.5fs and a minimum/maximum of 1.125ps/1.194ps respectively.

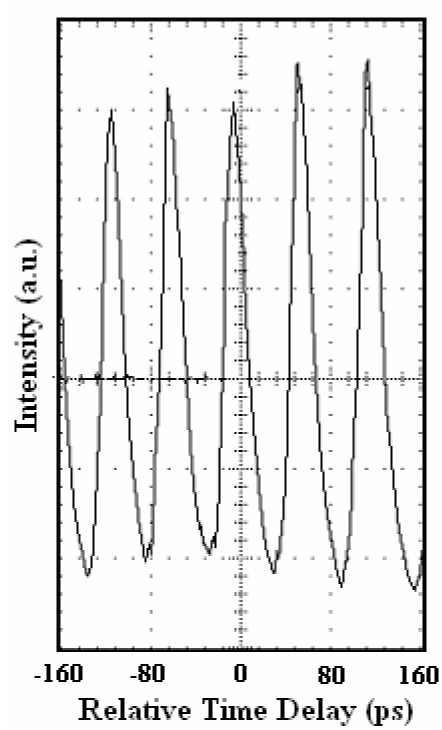

Figure 10. 18 GHz PMLS laser autocorrelator tracing. Depth of modulation was over 90%.

6. CONCLUSION

Though the theory and practical design aspects of PMLS lasers could fill a book, it is hoped that this paper will have addressed some of the theoretical and practical design aspects that have not been addressed previously. Because PMLS lasers have had problems with jitter and accuracy in past embodiments, not much research has been devoted to them for use in photonic clocking. It has only been over the last 5 years that research has been able to overcome some of their barriers to their commercial utilization such as power, jitter, and HVM. With our recent development of an 1100+nm material system for these lasers that is amenable to current HVM lines, the commercial viability of these lasers for microprocessor clocking has improved significantly. Not only do PMLS lasers offer a low-cost photonic alternative to current VCO technology but they also offers reduced jitter, skew, crosstalk, heat production, and power consumption while improving signal integrity.

Various PMLS laser characteristics and specifications can be tailored to the desired application through a variety of techniques. For instance, the PW can be decreased by either increasing the number of SAs, increasing the output facet's reflectivity, increasing the RB of the SA, decreasing the number of quantum wells, increasing the output frequency of the laser, or a combination thereof. Characteristics such as peak power, PW, jitter, and output frequency can be designed to meet various specifications. In regards to jitter, we found that our RMS jitter of 1ps was among the best for a PMLS laser. If even better jitter or accuracy specifications were needed then RF or optical injection at a (sub)harmonic would be an option.

As the viability of copper interconnects in microprocessors wanes and the demand for higher internal clock rates increases, photonic clocking will be inevitable. It is currently anticipated that HVM of microprocessors with photonic clocking will likely occur in the next 5-7 years. When that seismic change does occurs, the PMLS laser will be the most viable choice, not because of its scalability or other desirable characteristics, but because it is the most cost-effective solution.

REFERENCES

1. *The International Technology Roadmap for Semiconductors*, 2003 Edition.
2. J. Bautista, "The Potential Benefits of Photonics in the Computing Platform," *Optoelectronic Integrated Circuits VII*, Proceedings of SPIE Vol. 5729, pp. 1-8, 2005.
3. J. P. Siepmann, "Integration of a Photonic Clock Rate Multiplier with Global Distribution in Microprocessors," LightTime white paper available at http://www.lighttime.com/OPM-whitepaper.pdf, 2005.
4. W. W. Gibbs, "Computing at the Speed of Light," *Scientific American*, pp. 81-87, 2004.
5. C. Debaes, A. Bhatnagar, D. Agarwal, R. Chen, G. S. Keeler, N. C. Helman, H. Thienpont, and D. A. B. Miller, "Receiver-Less Optical Clock Injection for Clock Distribution Networks," *Selected Topics in Quantum Electronics, IEEE Journal of*, Volume **9**, Number 2, pp. 400-409, 2003.
6. A. V. Mule', E. N. Glytsis, T. K. Gaylord, and J. D. Meindl, "Electrical and Optical Clock Distribution Networks for Gigascale Microprocessors," *IEEE Transactions on Signal Very Large Scale Integration (VLSI) Systems*, Volume **10**, Number 5, pp. 582 – 594, 2004.
7. G. Hinton, D. Sager, M. Upton, D. Boggs, D. Carmean, A. Kyker, and P. Roussel, "The Microarchitecture of the Pentium 4 Processor," *Intel Technology Journal*, 2001.
8. J. A. Davis, R. Venkatesan, A. Kaloveros, M. Beylansky, S. J. Souri, K. Banerjee, K. C. Rahman, R. Reif, and J. D. Meindl, "Interconnect Limits on Gigascale Integration (GSI) in the 21st Century," *Proceedings of the IEEE*, Volume **89**, Number 3, pp. 305-324, 2001.
9. F. Iacona, D. Pacifici, A. Irrera, M. Miritello, G. Franzò, F. Priolo, D. Sanfilippo, G. Di Stefano, and P. G. Fallica, "Electroluminescence at 1.54 μm in Er-doped Si nanocluster-based devices," *Applied Physics Letters*, Volume **81**, Number 17, pp. 3242-3244.
10. J. R. Karin, R. J. Helkey, D. J. Derickson, R. Nagarajan, D. S. Allin, J. E. Bowers, and R. L. Thornton, "Ultrafast dynamics in field-enhanced saturable absorbers," *Applied Physics Letters*, Volume **64**, Number 6, pp. 676-678, 1994.
11. J. F. Martins-Filho, E. A. Avrutin, C. N. Ironside, and J. S. Roberts, "Monolithic Multiple Colliding Pulse Mode-Locked Quantum-Well Lasers: Experiment and Theory," *Selected Topics in Quantum Electronics, IEEE Journal of*, Volume **1**, Number 2, pp. 539-550, 1995.