

# A New Algorithm for Improved VDD Assignment in Low Power Dual VDD Systems

Sarvesh H Kulkarni

Ashish N Srivastava

Dennis Sylvester

University of Michigan, EECS Department, Ann Arbor, MI 48109

{shkulkar, ansrivas, dennis}@eecs.umich.edu

## ABSTRACT

We present the first in-depth study of the two existing algorithms, namely, Clustered Voltage Scaling (CVS) and Extended Clustered Voltage Scaling (ECVS), used for assigning the voltage supply to gates in integrated circuits having dual power supplies. We present a comparison of the achievable power savings using these algorithms on various benchmark circuits and first point out that ECVS does provide appreciably larger power improvements compared to CVS. We then provide a new algorithm based on ECVS that further improves the power savings by efficient assignment of the power supplies to the gates. Our new algorithm provides up to 66% power reduction and improves the power savings by up to 28% and 13% with respect to CVS and ECVS respectively. Furthermore, since level conversion is an essential component of dual power supply systems we also present the first circuit-specific sensitivity study of achievable power savings to the energy and delay penalties imposed by level conversion.

## Categories and Subject Descriptors

B.7.1 [Integrated Circuits]: VLSI.

## General Terms

Algorithms, Performance, Design, Experimentation

## Keywords

Dual VDD Design, Low Power Design Algorithms, CVS, ECVS, Level Converters

## 1. INTRODUCTION

Dynamic power dissipation in CMOS circuits is proportional to the square of the supply voltage (VDD). A reduction in VDD thus considerably lowers the power dissipation of the circuit. Dual (or more generally multi-) VDD design is an important scheme that exploits this concept to reduce power consumption in integrated circuits (ICs) [1, 2]. Since a reduction in VDD degrades circuit performance, in order to maintain performance in dual VDD designs, cells along critical paths are assigned to the higher power supply (VDDH) while cells along non-critical paths are assigned to a lower power supply (VDDL). Thus the timing slack available on non-critical paths is efficiently converted to energy

savings by use of a second supply voltage. However, level conversion (from VDDL to VDDH) becomes essential at boundaries where a VDDL driven cell drives a VDDH driven cell to eliminate the undesirable static current that otherwise flows. This current flows since the logic ‘HIGH’ signal of the VDDL driven cell cannot completely turn off the PMOS pull-up network of the subsequent VDDH cell.

The use of level converters is largely determined by the algorithm used in assigning VDD to gates. The two major existing algorithms used for VDD assignment are (1) Clustered Voltage Scaling (CVS) [1], and (2) Extended Clustered Voltage Scaling (ECVS) [3]. In CVS, the cells driven by each power supply are grouped (‘clustered’) together and level conversion is needed only at sequential element outputs (referred to as ‘*synchronous level conversion*’). In ECVS, the cell assignment is flexible, allowing level conversion anywhere (not just at the sequential element outputs) in the circuit. This is referred to as ‘*asynchronous level conversion*’. Since ECVS allows more freedom in VDD assignment, it has been suggested that it potentially provides greater power reductions than CVS [3]. However there has been no direct comparison between CVS and ECVS. In addition, nearly all subsequent work on dual-VDD synthesis has focused on CVS. In this work we demonstrate that ECVS is indeed far superior to CVS for a range of benchmark circuits in a 0.13 $\mu$ m CMOS technology.

Both CVS and ECVS assign the appropriate power supply to the gates by traversing the circuit from the primary outputs (PO) to the primary inputs (PI) in a leveled order. CVS is based on a topological constraint that allows a transition from a VDDH driven cell to a VDDL driven cell only, along any path from input to output (i.e., a VDDL driven cell may not feed into a VDDH driven cell). This greatly reduces the fraction of VDDL assigned gates and degrades the achievable power savings. Alternatively, ECVS relaxes this topological constraint by allowing a VDDL driven cell to feed a VDDH driven cell after undergoing level conversion through a dedicated asynchronous level converter (ALC). However, since ECVS performs this assignment simply by visiting gates one at a time in a reverse leveled manner, it still assigns supply voltages in a fundamentally constrained manner. Noting these drawbacks, we develop an algorithm that removes the ‘levelization’ approach to voltage assignment. Our algorithm is a greedy algorithm that assigns the power supply based on a sensitivity measure that considers slack changes as well as power savings *at each move* (the assignment of VDDL to a gate is referred to as a ‘*move*’ ahead in this paper), and hence results in superior power savings. It is related to the existing ECVS approach in that it allows the use of asynchronous level converters. Since level converters consume power and timing slack they have an important effect on the power savings. We

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED’04, August 9–11, 2004, Newport Beach, California, USA.

Copyright 2004 ACM 1-58113-929-2/04/0008...\$5.00.

investigate this impact and describe the importance of efficient level converters with respect to the total power savings.

The main contributions of this paper are:

- We conclusively demonstrate that ECVS is much more effective in reducing power compared to the more heavily studied CVS approach.
- We present a new algorithm that removes drawbacks in the existing ECVS algorithm to achieve uniformly lower total power.
- We study the impact of level converter performance overhead on total power dissipation over a range of benchmark circuits.

This paper is organized as follows. In Section 2, we briefly describe our implementations of CVS and ECVS, followed by Section 3, where we present a new enhanced algorithm. In Section 4, we describe the experimental setup. Section 5 presents the results obtained for various benchmark circuits using all three approaches. Finally, Section 6 concludes the paper.

## 2. BACKGROUND – CVS & ECVS

In this section we detail our implementations of the previously presented dual-VDD algorithms of CVS and ECVS. Since the original descriptions of these approaches were not highly detailed, this section seeks to provide a comprehensive summary of these algorithms. Note that for all the algorithms studied in this paper, we begin with a design with all cells assigned to VDDH and then assign VDDL according to the algorithm being applied. Both CVS and ECVS aim at utilizing the available timing slack in a circuit by applying a lower supply voltage on gates that are off the critical paths. This results in reduced dynamic power dissipation and hence lowers system level power dissipation. However, they differ in the policies they follow in making this power supply assignment. As a result of this, the final structure of the resultant netlists after applying these algorithms differs. As stated in Section 2, voltage level conversion becomes essential whenever a VDDL driven cell drives a VDDH driven cell. This is shown in Figure 1, where a VDDL driven inverter directly feeds into a VDDH driven inverter. The resulting DC current will result in extremely high static power dissipation without the use of level converters.




Figure 1. Demonstrating the need for level conversion

CVS and ECVS differ in the way they address the issue of level conversion. Since CVS does not allow for VDDL driven cells to directly feed VDDH driven cells, level conversion is therefore implemented only at flip-flop boundaries. This level conversion

functionality can be embedded into the flip-flop circuit [4, 5] and such flip-flops are hence called level converting flip-flops (LCFF).

ECVS relaxes this topological constraint and allows a VDDL driven cell to feed a VDDH driven cell after its output has undergone level conversion. ECVS thus has more freedom in finding parts of the circuit that can be operated at the lower supply and can potentially lead to higher power savings. However, the asynchronous level converters impose penalties in terms of their delay, power, and area. Fast and low power ALCs are thus important in mitigating these penalties. Figure 2 depicts the nature of the final topologies attained by CVS and ECVS when applied to a given circuit. From this figure, we note that CVS partitions a circuit into two clusters – one having only VDDH cells and the other having only VDDL cells. The scenario in which a VDDL driven cell directly feeds a VDDH driven cell is clearly precluded in this partitioning. On the other hand, ECVS allows interspersing of VDDL and VDDH cells with insertion of any required ALCs.




Figure 2. Resulting circuit structures after the application of CVS and ECVS.

We now discuss the implementation of these algorithms in greater detail. CVS maintains a list (referred to as  $L$ ) of candidate cells that can be assigned to VDDL. New cells continue to be added to this list as the algorithm proceeds. The elements of  $L$  are ranked according to a heuristic and the first element is chosen to be assigned to VDDL at each step of the algorithm. We follow the heuristics proposed in the first implementation of CVS [1] which ordered the cells in  $L$  on the basis of their load capacitance and slack.  $L$  is initialized to the set of gates that drive the circuit POs. Pseudo-code for CVS can be summarized as follows:

```

CVS () {
  STEP: 'SET VDDL' - Consider the first gate in  $L$  – for example 'A'.
    Set the supply voltage of A to VDDL.
    If A drives a PO, make the necessary LCFF insertion.
    Check timing.
    If this move is feasible (a 'feasible' move is one which allows the
    circuit to meet timing after it has been applied), {
      STEP: 'CONSTRAINED TOPOLOGY' -
        Add to  $L$ , all gates that fan into A but not into any other VDDH
        driven gate.
    }
    Else, {
      Reverse the move. Remove any added LCFFs.
    }
    Go back to STEP SET VDDL, if  $L$  is non-empty.
}

```

We note that the step '**CONSTRAINED TOPOLOGY**', guarantees that there will be no VDDL driven gate that feeds directly into a VDDH driven gate. However, this constraint acts to curtail many potential VDDL cell assignments as later results will demonstrate.

ECVS (as first implemented in [3]) begins by levelizing the circuit from the POs to the PIs. LCFFs are inserted when a cell driving a PO is assigned to VDDL (as in the case of CVS). Similarly ALCs are inserted whenever a VDDL gate feeds into a VDDH driven gate. The pseudo-code for the ECVS algorithm is as follows:

```

ECVS () {
  Traverse the circuit starting from the POs moving backwards
  according to the levels obtained initially.
  For each gate (ex: 'A') at each level, {
    STEP: 'SET VDDL' – Set the supply voltage of A to VDDL.
    For each gate B ∈ fanouts(A), {
      If (supply of B = VDDH)
        Insert an ALC on the path leading from A into B.
    }
    Check timing.
    Check power consumption.
    If this move is not feasible OR the power consumption
    increases by more than K%, {
      Reverse the move. Remove any added ALCs/LCFFs.
    }
  }
}

```

Since ECVS subsumes CVS, it can theoretically attain a much higher degree of VDDL gate assignments. ECVS however, must consider the overheads imposed by the ALCs. Also note that when a particular gate is assigned to VDDL the power dissipation can in fact increase. This can occur when a large number of ALCs must be inserted for a given move, or when the move itself leads to only a very small capacitance being set to VDDL. An increase in power dissipation due to a VDDL assignment is allowed (up to a threshold of K%), as an attempt to uncover better final solutions. The factor K therefore adds hill-climbing capability to the ECVS algorithm. This factor has an important effect on the finally achieved power savings as will be seen in Section 5 and hence must be chosen carefully.

Although ECVS has clear advantages over CVS, its policy of determining the VDD assignments is non-optimal and non-greedy. We address this in detail in the next section by describing an improved algorithm.

### 3. A NEW DUAL-VDD ASSIGNMENT ALGORITHM

The new algorithm described in this section is related to ECVS in that it allows the use of ALCs. We point out that ECVS-style approaches will be most effective if they are able to find “groups” or “clusters” of connected gates that can be assigned to the lower supply. This is true since such a grouped assignment will require fewer ALCs and minimize their resulting overhead. Note that the ALC delay penalty in particular should be avoided since it consumes valuable timing slack that could otherwise be translated into power savings. The total power of inserted ALCs should also be monitored since heavy use of ALCs will eventually lead to a rise in overall circuit power. In our algorithm, we introduce a sensitivity measure that uses the information available in the slack distribution of the circuit and the power savings attainable before finalizing *each VDDL assignment move*. This avoids the case of

ECVS which merely traverses the circuit (after levelization) and chooses the earliest *seen* feasible move.

At each stage of the algorithm, we calculate a sensitivity measure (to be defined shortly) for all cells that are potential candidates for VDDL assignment. Note that every VDDL assignment may call for either the insertion or removal of ALCs in the vicinity (at the inputs/output of the gate under consideration). This is because an ALC is required only when a VDDL driven gate needs to supply a VDDH driven gate. Note that since we do not follow a levelized VDDL assignment (as in ECVS), this ALC removal is frequently required and is accomplished by the `update_vicinity()` sub-routine in the pseudo-code below. As a result of a move, the arrival time at the output of the gate being assigned to VDDL will change (note that arrival time at the output includes the arrival time at the output of any added level converters, if the move required ALC/LCFF insertion). This changes the slack of various paths in the circuit. The overall power dissipation of the circuit will also change as a result of the move. A move assigning VDDL to a gate feeding a PO requires inserting an LCFF and in this case we must include the LCFF delay in the arrival time calculation. The LCFF data used in this work was obtained from [4, 5].

We define the sensitivity for a gate (say ‘x’) as follows:

$$\text{Sensitivity (gate 'x')} = \frac{\Delta \text{Power} \times \text{Slack at gate output}}{\Delta \text{Delay}}$$

where,  
 $\Delta \text{Power}$  = Change in total power due to move,  
 $\Delta \text{Delay}$  = Change in arrival time at gate output due to move.

Note that sensitivities for all gates that can undergo VDDL assignment are evaluated at every iteration of the algorithm and the move with the maximum sensitivity is selected. The state of the circuit is saved at this point and the algorithm proceeds to the next iteration. From the definition of sensitivity, we observe that this algorithm allows negative moves to be taken, thus opening the possibility of uncovering better solutions in the long run. Essentially this sensitivity measure enables us to choose the move giving *the best power savings per unit delay penalty*. The slack term in the sensitivity computation is used as a weighting factor to encourage VDDL assignment for gates with higher slack.

The new algorithm, designated GECVS (greedy-ECVS), tends to group VDDL gates together inherently due to the nature of the sensitivity function. Since the  $\Delta \text{Delay}$  and  $\Delta \text{Power}$  terms consider the ALC overheads associated with a particular VDDL assignment, the algorithm automatically guides itself towards building groups or clusters of VDDL gates. What is unique about GECVS is that these clusters can form *at the beginning of a path*, just as easily as they can at the end of a chain of combinational logic. This makes GECVS fundamentally more flexible than CVS or ECVS which proceed with VDDL assignment using a backwards traversal. Also note that for this as well as the preceding two algorithms, the best seen solution is saved at all times, which allows for the undoing of negative moves that did not lead to better final solutions. GECVS also differs from ECVS in the way the ALCs are placed. In GECVS, we place the ALCs directly at the input of the fanout gates rather than at the output of the driving gate in order to reduce the dynamic power consumed in switching the wire. This is particularly important given that a substantial portion of the total capacitance on a chip is due to

wiring and not devices [6]. The wireload model we use in our results is described in Section 4. Note that this slows the signal propagation speed on the interconnect (since the interconnect is now being driven by a VDDL gate) but reduces the overall dynamic power dissipation. This requirement can be easily incorporated in any placement tool.

Our algorithm can be summarized in pseudo-code form as follows:

```

GECVS ( )
{
  while (feasible moves remain) {
    Calculate sensitivities for all feasible moves.
    Choose the gate with the maximum sensitivity.
    Make the move.
    update_vicinity ( )
  }
}

```

## 4. SIMULATION SETUP

### 4.1 Process Technology

We use an industrial 0.13 $\mu$ m CMOS process for our study. The higher (nominal) power supply VDDH is 1.2V. Most earlier dual-VDD designs used a VDDL value of  $\sim 70\%$  of VDDH [7, 8, 9]. However, recent work [2, 10] has shown that the optimal value of VDDL for minimizing total power is  $\sim 50\%$  of VDDH. In this work, we hence analyzed the performance of the discussed algorithms for VDDL=0.6V [50% VDDH] and VDDL=0.8V [67% VDDH]. The threshold voltages for our technology are 0.12V and -0.09V for NMOS and PMOS respectively.<sup>1</sup> The wireload model we assumed was [11]:

$$\text{Wire Capacitance} = 5 \times [1 + (\text{number of fanouts} - 1) \times 0.4] \text{fF}$$

This translates to a wire capacitance of 7fF for a fanout of two, corresponding to a wirelength of approximately 35 $\mu$ m. Standard cell library files used in the experimental results for modeling delay and power were generated using comprehensive SPICE characterization of all gates used (including level converting gates).

### 4.2 Algorithm Implementation

The algorithms described in the preceding sections were implemented in C++ and tested on ISCAS85 [12] benchmarks, a 64-bit ALU, and a Huffman decoder. Note that we do not consider c17, c432, and c499 from the ISCAS85 benchmark set due to their small size. We also do not include c6288 as it is a completely balanced circuit and hence it was found to render all discussed dual VDD approaches equally ineffective in power reduction.

The synthesized design is initially sized using a sensitivity-based algorithm (similar to TILOS [13]) to obtain the complete power-delay curve for the design. The design is then resized from the initial synthesized point to a delay point that is backed off (i.e., the timing is relaxed somewhat) from the minimum achievable delay by a fixed percentage (this percentage is referred to as a

‘backoff’ point in the following sections). This approach is similar to the one used in [14]. Results for backoff points of 10% and 20% are reported in Section 5 (backoff points on the order of zero or a few percent are not viable operating points as they are in the steepest part of the circuit’s power-delay curve). The backoff point dictates the final cycle time of the circuit which is maintained during the application of CVS, ECVS, and GECVS.

## 4.3 Level Conversion

As described earlier, level conversion is an essential component of dual power supply systems. The overhead due to level conversion must be accounted for when reporting the achieved power savings.

As stated earlier, in the case of CVS, level conversion is required only at flip-flop boundaries (through the use of LCFFs). We have studied benchmarks that are purely combinational and thus incorporate the LCFF penalties by considering them to consume a fixed portion of the total delay budget (cycle time). We assume a LCFF delay penalty of 70ps [4, 5] for VDDL=0.6V/VDDH=1.2V. For VDDL=0.8V/VDDH=1.2V, the delay penalty is 50ps.<sup>2</sup>

In case of ECVS, ALCs are required in addition to LCFFs. The ALC we use is shown in Figure 3 [15]. This level converter has a delay of 70ps for a load of 17fF (representing a load of four drive 2X inverters and a typical wirelength of 35 $\mu$ m in our technology) for VDDL=0.8V/VDDH=1.2V. It consumes 22fJ of internal energy per transition. For VDDL=0.6V/VDDH=1.2V, the delay and energy are 95ps and 20fJ respectively.




Figure 3. High-speed asynchronous level converter.

## 5. RESULTS

### 5.1 Achievable Power Savings & Level Converter Analysis

Table 1 summarizes the dynamic power savings achieved by the described algorithms for the various benchmark circuits with VDDL = 0.6V and 0.8V. Note that we only compare the dynamic power savings since dual-VDD is targeted at this power component. Further, we do not consider dual-Vth processes that would be used in practice to address leakage power reduction. It is important to however note that we have ensured that the leakage power is a small fraction of the total power (16% for ALU64 and 4-8% for all remaining benchmarks, in the initial all VDDH design) in all our analyses. We assume an input switching

<sup>1</sup> We do not consider a dual-Vth process to simplify the reporting and interpretation of results.

<sup>2</sup> Note that the relevant LCFF delay penalty is the delay increase beyond a standard non-level-converting flip flop. A value of 50-70ps is on the order of 1.5-2 fanout-of-four inverter (FO4s) delays in the target technology.

activity of 0.1. The numbers reported in the tables below are dynamic power savings (percentages) with reference to the initial design (all cells at VDDH). Numbers are reported for two values of the backoff delay (20% in Table 1a and 10% in Table 1b).

**Table 1. Comparison of algorithm power savings.**  
(All numbers are percentage reductions relative to initial all VDDH design, except ALC column)

(a) Backoff = 20%

Benchmark	VDDL = 0.6V				VDDL = 0.8V			
	CVS	ECVS	GECVS	ALC	CVS	ECVS	GECVS	ALC
c880	24.1	28.4	35.2	9.8	20.8	24.2	27.3	8.0
c1908	7.3	8.6	13.8	7.2	6.9	7.0	12.1	5.6
c2670	20.6	28.9	32.1	9.2	16.3	23.3	26.3	7.9
c3540	4.2	23.0	30.7	9.3	3.3	18.9	24.6	5.4
c5315	27.4	35.9	39.3	10.8	22.1	28.0	31.8	8.2
c7552	33.9	39.8	44.2	12.3	26.4	31.2	35.4	9.6
ALU64	56.2	62.9	65.7	12.7	46.2	49.5	51.5	9.1
HUFFMAN	18.7	19.9	30.2	14.8	14.9	17.6	25.2	12.5
AVERAGE	24.0	30.9	36.4	10.8	19.6	25.0	29.3	8.3

(b) Backoff = 10%

Benchmark	VDDL = 0.6V				VDDL = 0.8V			
	CVS	ECVS	GECVS	ALC	CVS	ECVS	GECVS	ALC
c880	21.1	26.3	30.6	9.2	17.2	20.3	24.8	6.5
c1908	5.8	6.9	10.7	5.2	4.6	5.4	8.8	6.0
c2670	16.7	24.4	26.9	5.7	13.1	18.1	23.5	7.9
c3540	3.5	14.0	24.3	8.3	2.9	12.3	20.8	7.1
c5315	22.7	31.8	35.5	10.3	18.9	25.3	28.4	8.2
c7552	30.4	35.5	39.4	10.5	24.1	27.5	31.0	9.0
ALU64	54.6	61.2	63.6	13.2	45.0	48.1	49.9	8.8
HUFFMAN	11.9	16.3	25.6	13.5	10.6	15.5	21.0	11.9
AVERAGE	20.8	27.1	32.1	9.5	17.0	21.6	26.0	8.2

We observe that the proposed algorithm outperforms both existing algorithms for all studied benchmarks, values of VDDL, and delay backoff. Also, both ECVS and GECVS perform significantly (in some circuits providing approximately twice the power savings or more) better than CVS because of the greater VDDL assignment as demonstrated in the next sub-section. On average, GECVS optimized circuits show 6-8% (11-16%) lower power than with ECVS (CVS) for the various VDDL and backoffs considered. Some circuits exhibit much larger power savings; in particular the application of GECVS to the Huffman decoder and c3540 (an 8-bit ALU) yields up to 13% (28%) lower power than ECVS (CVS). In general, GECVS performs best relative to the other algorithms at the 20% backoff point and VDDL=0.6V; these conditions (i.e., an aggressively scaled VDDL which is about half of VDDH and a correspondingly scaled Vth) also lead to the maximum overall power savings [10], indicating that GECVS is ideally suited to scenarios in which dual VDD would be a good option.


We also report the percentage of power consumed by ALCs in the second column under GECVS in Table 1. Our findings that ALCs consume 8-10% of total power on average across benchmarks are in keeping with [3] which claimed a level converter power overhead of 8%.

In this study, all algorithms discussed achieve power savings by simply changing the supply of the gates in the design to the lower supply (VDDL) whenever possible. It is important to note that gate resizing and dual-Vth techniques can be concurrently

employed to further reduce the power dissipation, motivating further work combining these three optimization tools [14].

## 5.2 Standard Cell Statistics

Figure 4 compares the fraction of the total gates assigned to VDDL in the final design after the application of the three described algorithms (at VDDL = 0.8V, backoff = 10%). We note that the new algorithm achieves 21% and 7% higher VDDL assignment compared to CVS and ECVS respectively (averaged over all benchmarks). This improvement directly translates into the higher power savings reported above.



**Figure 4. Comparison of achieved levels of VDDL assignment by CVS/ECVS/GECVS.**

## 5.3 Power Sensitivity to Level Converter Performance

We studied the effect that level converter performance has on the total power savings achieved by the proposed algorithm in several benchmarks. The level converter shown in Figure 3 improves delay by up to 20% (for fixed energy) or reduces power dissipation by up to 50% (for fixed delay) compared to previously existing level converters [15]. By scaling the characterized data for this level converter by these scaling factors, we obtained results that reflect the impact that inferior level converters have on system-level power dissipation. For c880, c3540, and c7552, the level converter we used provided a 5.9%, 9.0% and 5.9% increase in power savings by virtue of its 50% lower power consumption or alternatively a much smaller 0.2%, 1.4% and 0.2% increase in power savings due to the potential for 20% faster operation.

Figure 5 shows the variation in the achieved power savings for the Huffman benchmark as ALC delay and power are varied (holding one fixed at a time; VDDL=0.6V, backoff = 10%). We can see a 9.5% increase in power savings for a 50% energy reduction and a 3.2% increase in power saving by virtue of the 20% speed up. The system-level power savings shows a higher sensitivity to delay for this benchmark as compared to c880, c3540 and c7552 (at the same VDDL and backoff conditions). This may be due to a higher number of moderately sub-critical paths on which VDDL gates can just be inserted (i.e., if the ALC delay is increased slightly, these paths can no longer accept VDDL assignments). However, it is difficult to observe this level of detail in the path delay distribution of any reasonably sized circuit. Similar trends to those described in this section were observed for the other benchmarks studied in this paper.




Figure 5. Impact of ALC performance on system level power dissipation.

### 5.4 Impact of Hill-Climbing Factor (K) on ECVS Power

It was noted in Section 2 that the factor K has an important effect on the power savings provided by the ECVS algorithm. Figure 6 demonstrates this effect. A strong dependency of the power savings on K can be observed. Note that in reporting the power savings of ECVS in Section 5.1, we use a value of  $K=0.1\%$  which we have found to give the best results. We therefore report the best savings possible by ECVS in Section 5 when making comparisons with CVS and GECVS.




Figure 6. Impact of K on ECVS power savings.

## 6. CONCLUSIONS

In this paper, we demonstrated that the ECVS approach does indeed provide much improved power savings beyond the more traditional, and well-studied, CVS approach. We then proposed a new algorithm that enables more flexible VDDL assignment and presented a comparison of this new GECVS algorithm to CVS and ECVS. The proposed algorithm shows a 7% (13%) average and 13% (28%) maximum improvement in total power savings over ECVS (CVS). We finally showed that level converters have a significant impact on the system-level power dissipation, motivating further work in the development of fast and low-energy asynchronous level converters.

## 7. ACKNOWLEDGMENTS

This work was supported by the Semiconductor Research Corporation, the MARCO/DARPA Gigascale Systems Research Center, and Intel Corporation. The authors thank Professor

Borivoje Nikolic from the University of California, Berkeley for providing the Huffman benchmark.

## 8. REFERENCES

- [1] K. Usami and M. Horowitz, "Clustered voltage scaling technique for low-power design," *Proc. ISLPED*, pp. 3-8, 1995.
- [2] C. Chen, A. Srivastava and M. Sarrafzadeh, "On gate level power optimization using dual-supply voltages," *IEEE Trans. on VLSI Systems*, vol. 9, pp. 616-629, Oct. 2001.
- [3] K. Usami, M. Igarashi, F. Minami, M. Ishikawa, M. Ichida and K. Nogami, "Automated low-power technique exploiting multiple supply voltages applied to a media processor," *IEEE JSSC*, pp. 463-472, Mar. 1998.
- [4] M.R. Bai and D. Sylvester, "Analysis and design of level converting flip-flops for dual-Vdd/Vth integrated circuits," *Proc. International Symposium on System-on-Chip*, pp. 151-154, 2003.
- [5] F. Ishihara, F. Sheikh and B. Nikolic, "Level conversion for dual supply systems," *Proc. ISLPED*, pp. 164-167, 2003.
- [6] N. Magen, A. Kolodny, U. Weiser and N. Shamir, "Interconnect power dissipation in a microprocessor," *International Workshop on System-Level Interconnect Prediction*, pp. 7-13, 2004.
- [7] M. Hamada, Y. Ootaguro and T. Kuroda, "Utilizing surplus timing for power reduction," *Proc. CICC*, pp. 89-92, 2001.
- [8] M. Takahashi, M. Hamada, T. Nishikawa, H. Arakida, T. Fujita, F. Hatori, *et al.*, "A 60-mW MPEG4 video codec using clustered voltage scaling with variable supply-voltage scheme," *IEEE JSSC*, pp. 1772-1780, Nov. 1998.
- [9] T. Kuroda and M. Hamada, "Low-power CMOS digital design with dual embedded adaptive power supplies," *IEEE JSSC*, pp. 652-655, Apr. 2000.
- [10] A. Srivastava and D. Sylvester, "Minimizing total power by simultaneous Vdd/Vth assignment," *Proc. ASPDAC*, pp. 400-403, 2003.
- [11] D. Sylvester and K. Keutzer, "System-level performance modeling with BACPAC - Berkeley advanced chip performance calculator," *International Workshop on System-Level Interconnect Prediction (workshop notes)*, pp. 109-114, 1999.
- [12] F. Brglez and H. Fujiwara, "A neural netlist of 10 combinational benchmark circuits and a target translator in Fortran," *Proc. ISCAS*, pp. 695-698, May 1985.
- [13] J. Fishburn and A. Dunlop, "TILOS: a posynomial programming approach to transistor sizing," *Proc. ICCAD*, pp. 326-328, 1985.
- [14] A. Srivastava, D. Sylvester and D. Blaauw, "Power minimization using simultaneous gate sizing, dual-Vdd, and dual-Vth assignment," *Proc. DAC*, in press, 2004.
- [15] S. H. Kulkarni and D. Sylvester, "High performance level conversion for dual VDD design," *IEEE Trans. On VLSI Systems*, in press, 2004.