

A Low-Power Bus Design Using Joint Repeater Insertion and Coding

Srinivasa R. Sridhara and Naresh R. Shanbhag
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
1308 W Main St., Urbana IL 61801
[sridhara,shanbhag]@uiuc.edu

ABSTRACT

In this paper, we propose joint repeater insertion and crosstalk avoidance coding as a low-power alternative to repeater insertion for global bus design in nanometer technologies. We develop a methodology to calculate the repeater size and separation that minimize the total power dissipation for joint repeater insertion and coding for a specific delay target. This methodology is employed to obtain power vs. delay trade-offs for 130-nm, 90-nm, 65-nm, and 45-nm technology nodes. Using ITRS technology scaling data, we show that proposed technique provides 54%, 67%, and 69% power savings over optimally repeater-inserted 10-nm 32-bit bus at 90-nm, 65-nm, and 45-nm technology nodes, respectively, while achieving the same delay.

Categories and Subject Descriptors: B.4.3 [Input/output and data communications]: Interconnections (Subsystems)

General Terms: Design, Performance

Keywords: Crosstalk, low-power, repeaters, coding

1. INTRODUCTION

Low-power and high-performance operation is necessary for all components in microprocessors and system-on-chip (SOC) designs. This is especially true for global buses, whose delay and power dissipation show an increasing trend with technology scaling. According to International Technology Roadmap of Semiconductors (ITRS) [1] gate delay reduces with scaling, while global wire delay increases. Therefore, delay of global buses will act as the performance bottleneck in many high-performance system-on-chip (SOC) designs. Further, interconnection networks consume 20%-36% of total system power in many large SOCs [2]. Future SOCs are expected to follow the network-on-chip (NOC) paradigm [3], where high-speed energy-efficient communication between various SOC components is vital.

Repeater insertion [4] is able to reduce the growing gap between logic and interconnect delay. However, repeater insertion increases the power dissipation of bus due to large buffers needed to drive the bus. In nanometer technologies, leakage power dissipation in these buffers can account for more than 20% of total power consumption [5].

Crosstalk avoidance coding (CAC) [6-10], has emerged

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED'05, August 8-10, 2005, San Diego, California, USA
Copyright 2005 ACM 1-59593-137-6/05/0008 ...\$5.00.

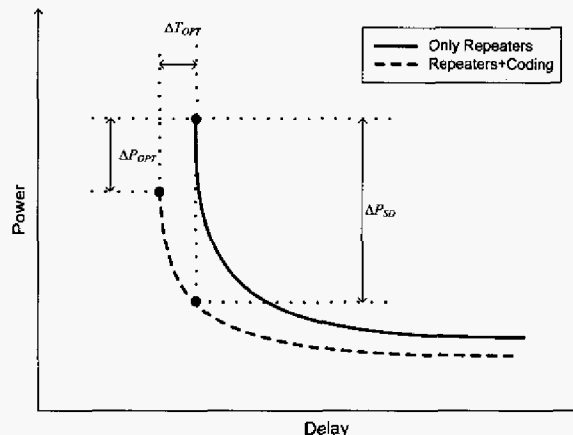


Figure 1: Power vs. delay trade-offs achieved using repeater insertion and joint repeater insertion and coding.

as an attractive technique for reducing delay and power. CAC encodes the input data to eliminate the worst-case coupling transitions in a bus and reduces the worst-case delay. Elimination of worst-case transitions also reduces average transition activity and, hence, reduces power dissipation. Codes have been proposed that reduce the bus power dissipation by more than 20% [9]. However, the best possible delay reduction via CAC is limited to 75% for buses without repeaters [6] and is much lower for buses with repeaters. Clearly, CAC cannot achieve the delay reduction that is achievable by repeater insertion, which achieves more than one order of magnitude reduction for nodes beyond the 65-nm technology [1].

In this paper, we propose joint repeater insertion and coding as a low-power alternative to repeater insertion for global buses. Joint repeater insertion and coding combines the delay reduction benefits of repeater insertion with delay and power reduction benefits of coding. The proposed technique achieves greater delay reduction than repeater insertion alone, while consuming lower power. First, we develop a methodology to calculate the repeater size and separation that minimize the total power dissipation for joint repeater insertion and coding for a specific delay target. This methodology extends the methodology proposed in [5] by including the effects of coding and codec overhead on total delay and power dissipation. The methodology employs a bus power model that includes dynamic, short-circuit, and leakage power dissipation for the repeaters and the coders, and dynamic power dissipation for the interconnects.

Next, we employ the methodology to obtain power vs. delay trade-offs in global bus design for 130-nm, 90-nm, 65-

nm, and 45-nm technology nodes. An illustration of power vs. delay trade-off achieved using the methodology is shown in Fig. 1. The figure shows that the proposed technique has lower power dissipation than repeater insertion for all delay values. In Section 4, we show that this is indeed the case for 65-nm and 45-nm technology nodes. Fig. 1 also shows ΔP_{OPT} , ΔT_{OPT} , and ΔP_{SD} as metrics employed in comparing the proposed technique with repeater insertion. Here, ΔP_{OPT} denotes the power savings and ΔT_{OPT} denotes the delay reduction. Further, ΔP_{SD} denotes the power savings achieved by the proposed scheme while achieving the same delay as optimal repeater insertion as shown in the figure. Using ITRS technology scaling data, we show that joint repeater insertion and coding provides ΔP_{SD} of 54%, 67%, and 69% power savings for a 10-mm 32-bit bus at 90-nm, 65-nm, and 45-nm technology nodes, respectively. At 45-nm node, we show that delay-optimal joint repeater insertion and coding achieves $\Delta T_{OPT} = 40\%$ and $\Delta P_{OPT} = 21\%$ over delay-optimal repeater insertion.

2. JOINT REPEATER INSERTION AND CODING

Consider a joint repeater insertion and coding scheme for a k -bit bus. The k input bits are encoded into n coded bits at the transmitter by employing a crosstalk avoidance code. The n coded bits are transmitted over n parallel wires by employing drivers of size s . After every segment of length l , n repeaters of size s are inserted. At the receiver, the received bits are decoded back into the original k data bits. The total length of the bus is L and the interconnects have resistance r per unit length, bulk capacitance c_b per unit length, and coupling capacitance c_c per unit length.

2.1 Uncoded repeater-inserted bus

The effective interconnect capacitance seen by the repeaters depends on the transitions occurring on its adjacent wires due to crosstalk. This capacitance can be modeled as [6]

$$c_a = c_b + pc_c, \quad (1)$$

where p is referred to as the coupling factor. The coupling factor takes values $p = 0, 1, 2, 3$, and 4 depending on the transitions occurring on adjacent wires [6]. The delay of a wire in the bus with crosstalk is obtained by modifying the expression in [4] as follows

$$T_{wire}(p) = L \log_e 2 \left(\frac{1}{l} \tau_s (c_i + c_o) + \frac{\tau_s}{s} (c_b + pc_c) + rsc_i + \frac{1}{2} r (c_b + pc_c) l \right), \quad (2)$$

where τ_s , c_i , and c_o are output resistance, input capacitance, and output capacitance of a minimum sized inverter, respectively.

When coding is not employed, p can take the worst-case value of 4. Therefore, the worst-case delay of an uncoded repeater-inserted bus is given by

$$T_{rep} = T_{wire}(4), \quad (3)$$

The power consumption of an uncoded repeater-inserted bus is composed of dynamic, short-circuit, and leakage power dissipation in the repeaters and dynamic power dissipation in the interconnects. We express the total power as a function of self transition activity α and coupling transition activity β by modifying the expression in [11]. The total power P_{rep} is given by

$$P_{rep}(\alpha, \beta) = kL \left(k_1 \left(\alpha \left(\frac{s}{l} (c_o + c_i) + c_b \right) + 2\beta c_c \right) + k_2 \frac{s}{l} + k_3 \alpha s \frac{T_{wire}(p)}{L \log_e 2} \right), \quad (4)$$

where

$$\begin{aligned} k_1 &= V_{DD}^2 f_{clk} \\ k_2 &= \frac{1}{2} V_{DD} (I_{off_n} + 2I_{off_p}) W_{n_{min}} \\ k_3 &= V_{DD} W_{n_{min}} I_{sc} f_{clk} \log_e 3. \end{aligned}$$

Here, V_{DD} is the power supply voltage, f_{clk} is the clock frequency, I_{off_n} (I_{off_p}) is the leakage current per unit NMOS (PMOS), $W_{n_{min}}$ is the width of the NMOS transistor in minimum sized inverter, and I_{sc} is the per unit width short-circuit current.

2.2 Effect of coding

CAC modifies delay and power dissipation in the following manner. CAC reduces the maximum coupling factor from $p = 4$ to $p = 1, 2$, or 3 depending on the choice of the code. For example, the forbidden transition overlapping code (FTOC) [9] has a coupling factor of $p = 2$, while the one lambda code (OLC) [10] has $p = 1$. The reduction in effective interconnect capacitance reduces the worst-case delay of the bus. However, coding adds latency in the form of encoder and decoder delays. Therefore, the worst-case delay of the bus from the input of the transmitter to the output of the receiver is given by

$$T_{joint} = T_{wire}(p) + T_{codec}, \quad (5)$$

where T_{codec} is the delay of the codec (encoder and decoder).

Crosstalk avoidance coding modifies the transition activity on the bus. It has been shown in [9, 10] that CAC reduces total coupling transition activity $\frac{n}{k}\beta$, but increases the self transition activity $\frac{n}{k}\alpha$. There is an effective reduction in power dissipation if the coupling capacitance dominates the bulk capacitance, as is the case in nanometer technologies. The total power dissipation of the joint scheme is given by

$$P_{joint} = \frac{n}{k} P_{rep}(\hat{\alpha}, \hat{\beta}) + P_{codec}, \quad (6)$$

where $\hat{\alpha}$ and $\hat{\beta}$ are self and coupling transition activities with coding and P_{codec} is the power consumption of the codec.

3. METHODOLOGY

As described in Section 2, both delay and power dissipation are functions of segment length l and repeater size s . It is well known that delay optimal values of l and s exist [4]. However, these optimally sized and separated repeaters dissipate large amounts of power [5]. In many instances, such as non-critical global buses, a target delay is desired rather than minimal delay. In such cases, it is possible to reduce power dissipation by employing smaller repeaters and larger segment lengths. Further, it has been shown in [5] that the delay is very shallow with respect to both repeater size and segment length close to the minimum point. Therefore, large savings in power dissipation are possible by tolerating a small delay penalty. Here, we outline a methodology for power optimization while satisfying a given constraint on delay. It is derived by extending the methodology in [5] to include coupling factor p , coupling activity β , codec delay T_{codec} , and codec power dissipation P_{codec} .

If the desired delay is T_{target} , then we set

$$T_{wire}(p) + T_{codec} \leq T_{target}. \quad (7)$$

If T_{target} is greater than the optimal delay, then there are several combinations of l and s that satisfy (7). In order to optimize for power, we determine the values of l and s that minimize (6), subject to the constraint in (7). This is done by setting the partial derivative of P_{joint} in (6) with respect to s to zero. This exercise, along with (7), leads to

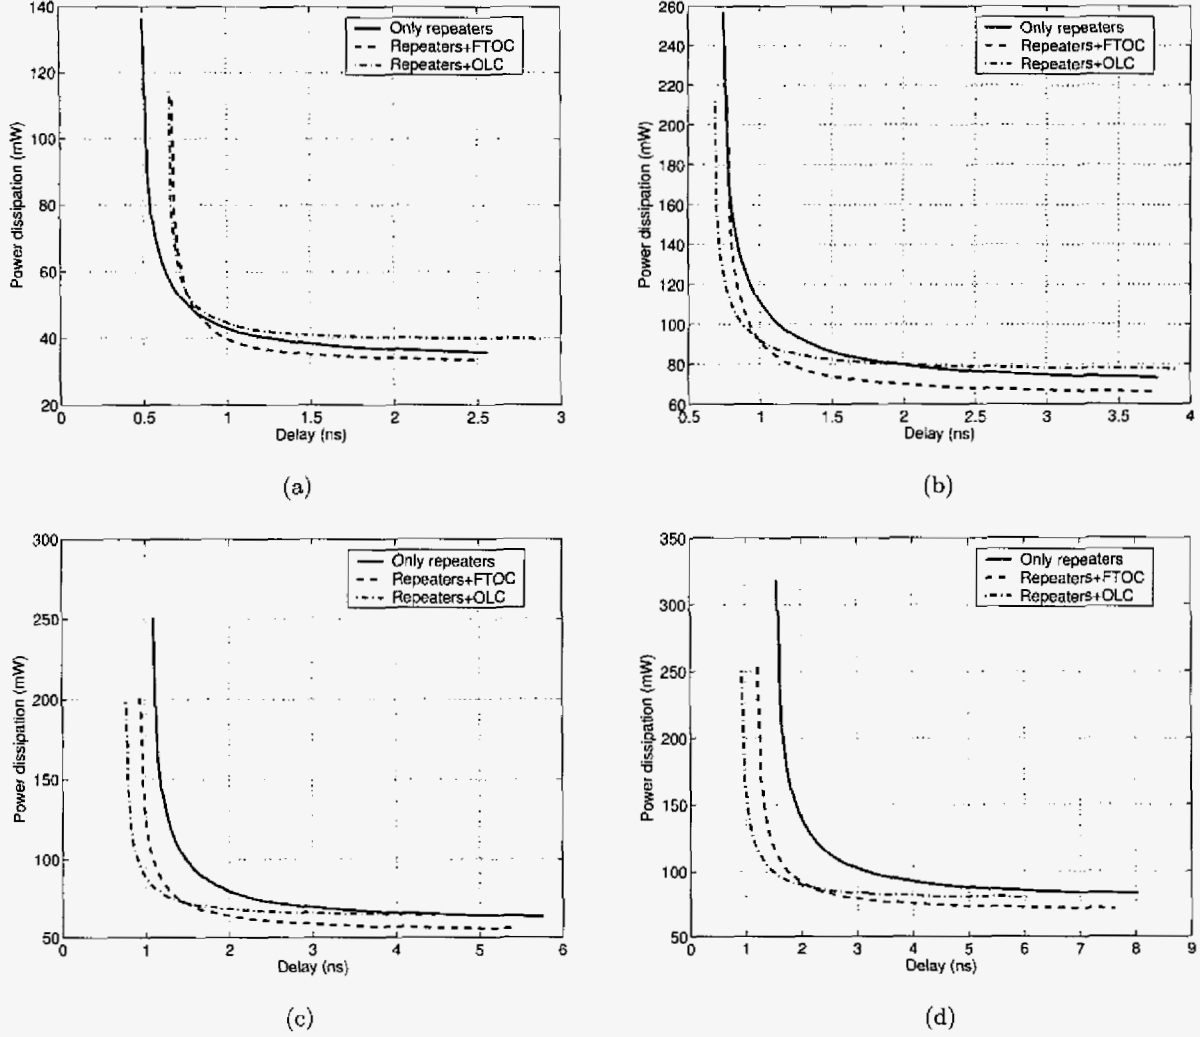


Figure 2: Power vs. delay trade-off for 10-mm 32-bit bus: (a) 130 nm, (b) 90 nm, (c) 65 nm, and (d) 45 nm.

the following three non-linear equations [5]:

$$\begin{aligned} \frac{k_1 \hat{\alpha} (c_i + c_o)}{l} + \frac{k_2}{l} + k_3 \hat{\alpha} \frac{T_{target} - T_{codec}}{L \log_e 2} \\ - \left[\frac{k_1 s (c_i + c_o) + k_2 s}{l^2} \right] \frac{dl}{ds} = 0 \\ \frac{1}{l} r_s (c_i + c_o) + \frac{r_s}{s} (c_b + p c_c) + r s c_i \\ + \frac{1}{2} r (c_b + p c_c) l - \frac{T_{target} - T_{codec}}{L \log_e 2} = 0 \quad (8) \\ \left[\frac{1}{2} r (c_b + p c_c) - \frac{r_s (c_i + c_o)}{l^2} \right] \frac{dl}{ds} + r c_i \\ - \frac{r_s (c_b + p c_c)}{s^2} = 0. \end{aligned}$$

The above equations are solved numerically using Newton-Raphson to obtain the values of s and l .

4. RESULTS

The methodology described in Section 3 is used to design a 32-bit global bus in the top metal layer for various technology nodes. The ITRS technology parameters are shown in Table 1. The interconnect capacitance values were obtained using RLC modeling in HSPICE. Device parameters have been obtained from [11]. We assume that the input data is spatially and temporally uncorrelated with "0" and "1" being equiprobable. Codec overhead is estimated from synthesized gate-level netlists obtained using a 130-nm CMOS standard cell library. The ITRS scaling trend for delay and power dissipation is employed to estimate the overheads for other technology nodes. Due to the codec overhead, the effective reduction in delay and power dissipation with coding will depend on bus length L . In this paper, we consider the design of a 10-mm global bus.

Fig. 2 shows the power dissipation for for uncoded and coded buses with repeater insertion for the four technology nodes: 130-nm, 90-nm, 65-nm, and 45-nm. Each curve represents a power vs. delay trade-off achieved by employing a specific joint scheme in the design of a 10-mm 32-bit bus. The left most point of each curve represents the delay-optimized solution and, hence, consumes the highest power.

Table 1: Technology and device parameters for various technology nodes based on the ITRS [11]. The interconnect data is for top metal layer.

Node (nm)	130	90	65	45
width (nm)	335	230	145	103
thickness (nm)	670	482	319	236
dielectric (μm)	6.3	4.7	3.9	2.9
c_b (fF/mm)	11.7	11.8	9.4	10.0
c_c (fF/mm)	90.4	93.3	96.7	99.9
r_s (k Ω)	6.23	9.04	9.6	13.2
c_i (fF)	1.33	1.1	1.03	0.9
c_o (fF)	3.32	2.04	1.22	0.6
V_{DD} (V)	1.1	1	0.7	0.6
I_{offn} ($\mu\text{A}/\mu\text{m}$)	2	3.56	20	35.5
I_{offp} ($\mu\text{A}/\mu\text{m}$)	1.34	2.38	13.4	23.83
I_{sc} ($\mu\text{A}/\mu\text{m}$)	65	65	65	65
f_{clk} (GHz)	1.68	3.99	6.73	11.51

Table 2: Benefits of joint repeater insertion and coding over delay-optimized uncoded 10-mm 32-bit bus.

Scheme	Metric	130	90	65	45
Repeaters + FTOC	ΔP_{SD} (%)	-	-	60.9	65.9
	ΔP_{OPT} (%)	18.1	18.6	19.9	20.5
	ΔT_{OPT} (%)	-35.0	-3.50	13.7	22.2
Repeaters + OLC	ΔP_{SD} (%)	-	53.8	67.0	69.4
	ΔP_{OPT} (%)	16.3	17.5	21.0	21.4
	ΔT_{OPT} (%)	-31.4	8.0	29.8	40.4

While the rest of points on the curve show the power dissipation for a design with the given delay penalty.

At the 130 nm node in Fig. 2(a), it is seen that the optimal delay for both FTOC and OLC joint schemes is greater than the optimal delay for the uncoded case. Therefore, we conclude that coding does not provide any delay benefits at the 130 nm node due to the high cost of encoding and decoding. However, as we move along technology scaling path to 90 nm, we observe that the power dissipation and delay of the uncoded bus increases more rapidly than the coded bus. This is because the delay and power dissipation of interconnects increase with scaling, while the delay and power dissipation of the coders reduce. In Fig. 2(b), we observe that the OLC-coded bus provides significant delay and power reduction near the delay-optimized solution point.

Employing the metrics defined in Fig. 1, the OLC-coded bus provides power savings of $\Delta P_{SD} = 54\%$ while achieving the same delay as the delay-optimized repeater bus. The achieved savings improve with scaling as seen in Fig. 2(c) and Fig. 2(d). Table 2 lists the three metrics computed by using the plots in Fig. 2. There is no data for ΔP_{SD} at the 130-nm node as joint repeater insertion and coding has total delay greater than the uncoded delay-optimized bus. However, joint repeater insertion and OLC provides ΔP_{SD} of 54%, 67%, and 69% at 90-nm, 65-nm, and 45-nm nodes, respectively. Joint repeater-insertion and FTOC also provides ΔP_{SD} of 61% and 66% for 65-nm and 45-nm technologies, respectively.

In addition, it is seen that the trade-off curves of joint repeater insertion and coding are always below the curves of the uncoded bus for 65-nm and 45-nm nodes. This indicates that OLC and FTOC provides power savings over uncoded bus at all target delay values in 65-nm and 45-nm nodes.

When the optimal delay design of joint repeater insertion and coding is compared with uncoded repeater insertion for 10-mm 32-bit bus, we see that joint repeater in-

sertion and FTOC has negative ΔT_{OPT} for 130-nm and 90-nm nodes indicating that the codec overhead nullifies the effect of crosstalk avoidance at these nodes. However, FTOC achieves ΔT_{OPT} of 14% and 22% for 65-nm and 45-nm nodes, respectively. Similarly, OLC achieves ΔT_{OPT} of 8%, 30%, and 40% for 90-nm, 65-nm, and 45-nm nodes, respectively. These delay gains are accompanied by power savings ΔP_{OPT} as shown in the table. At 45-nm node, joint repeater insertion and OLC achieves $\Delta P_{OPT} = 21\%$ along with $\Delta T_{OPT} = 40\%$ compared to optimally repeater-inserted uncoded bus.

5. CONCLUSIONS

We have shown that joint repeater insertion and coding provides the best power-delay trade-off in 90-nm and smaller technologies. We have shown that the proposed scheme has lower delay than repeater insertion, while consuming lower amounts of power. We have shown that joint scheme achieves significant power reduction over optimally repeatered bus, while achieving the same delay.

While we have quantified the delay and power benefits of joint repeater insertion and coding, another metric that is hard to quantify is the effect on area and routing congestion. Repeater insertion causes routing congestion in lower metal layers. Joint scheme is able to achieve the same delay as repeater insertion with larger separation between repeaters. This alleviates some of the routing congestion in the lower metal layers. Further, joint schemes requires smaller repeaters and, hence, requires lower silicon area. However, these area savings and reduction in routing congestion are achieved at the cost of increased area due to additional wires and the codec. Whether this trade-off between routing congestion at the lower metal layers and additional area at the global bus layer is beneficial for a given design depends on the specific design and is a subject of future research.

6. ACKNOWLEDGMENTS

This work was supported by the MARCO-sponsored Gascale Systems Research Center and Intel Corporation.

7. REFERENCES

- [1] *International Technology Roadmap for Semiconductors*, Semiconductor Industry Association, 2003. [Online]. Available: <http://public.itrs.net>
- [2] V. Soteriou and L.-S. Peh, "Design-space exploration of power-aware on/off interconnection networks," in *Proc. ICCD*, 2004, pp. 510-517.
- [3] L. Benini and G. D. Micheli, "Networks on chips: a new SoC paradigm," *IEEE Computer*, vol. 35, pp. 70-78, Jan. 2002.
- [4] H. Bakoglu, *Circuits, Interconnects and Packing for VLSI*. Reading, MA: Addison-Wesley, 1990.
- [5] K. Banerjee and A. Mehrotra, "A power-optimal repeater insertion methodology for global interconnects in nanometer designs," *IEEE Trans. Electron Devices*, vol. 49, pp. 2001-2007, Nov. 2002.
- [6] P. P. Sotiriadis and A. Chandrakasan, "Reducing bus delay in submicron technology using coding," in *Proc. ASP-DAC*, 2001, pp. 109-114.
- [7] C. Duan, A. Tirumala, and S. P. Khatri, "Analysis and avoidance of cross-talk in on-chip buses," in *Proc. Hot Interconnects*, 2001, pp. 133-138.
- [8] B. Victor and K. Keutzer, "Bus encoding to prevent crosstalk delay," in *Proc. ICCAD*, 2001, pp. 57-63.
- [9] S. R. Sridhara, A. Ahmed, and N. R. Shanbhag, "Area and energy-efficient crosstalk avoidance codes for on-chip buses," in *Proc. ICCD*, 2004, pp. 12-17.
- [10] S. R. Sridhara and N. R. Shanbhag, "Coding for reliable on-chip buses: fundamental limits and practical codes," in *Proc. VLSI Design*, 2005, pp. 417-422.
- [11] M. L. Mui, K. Banerjee, and A. Mehrotra, "A global interconnect optimization scheme for nanometer scale vlsi with implications for latency, bandwidth, and power dissipation," *IEEE Trans. Electron Devices*, vol. 51, pp. 195-203, Feb. 2004.