WATTLE: A Trainable Gain Analogue VLSI Neural Network

Richard Coggins and Marwan Jabri
Systems Engineering and Design Automation Laboratory
Department of Electrical Engineering J03,
University of Sydney, 2006.
Australia.
Email: richardc@sedal.su.oz.au
marwan@sedal.su.oz.au

Abstract

This paper describes a low power analogue VLSI neural network called Wattle. Wattle is a 10:6:4 three layer perceptron with multiplying DAC synapses and on chip switched capacitor neurons fabricated in 1.2um CMOS. The on chip neurons facillitate variable gain per neuron and lower energy/connection than for previous designs. The intended application of this chip is Intra Cardiac Electrogram classification as part of an implantable pacemaker/defibrillator system. Measurements of the chip indicate that 10pJ per connection is achievable as part of an integrated system. Wattle has been successfully trained in loop on parity 4 and ICEG morphology classification problems.

1 INTRODUCTION

A three layer analogue VLSI perceptron has been previously developed by [Leong and Jabri, 1993]. This chip named Kakadu uses 6 bit digital weight storage, multiplying DACs in the synapses and fixed value off chip resistive neurons. The chip described in this paper called Wattle has the same synapse arrays as Kakadu, however, has the neurons implemented as switched capacitors on chip. For both Kakadu and Wattle, analogue techniques have been favoured as they offer greater opportunity to achieve a low energy and small area design over standard digital

Figure 1: Wattle Synapse Circuit Diagram

techniques since the transistor count for the synapse can be much lower and the circuits may be biased in subthreshold. Some work has been done in the low energy digital area using subthreshold and optimised threshold techniques, however no large scale circuits have been reported so far. [Burr and Peterson, 1991] The cost of using analogue techniques is however, increased design complexity, sensitivity to noise, offsets and component tolerances. In this paper we demonstrate that difficult nonlinear problems and real world problems can be trained despite these effects.

At present, commercially available pacemakers and defibrillators use timing decision trees implemented on CMOS microprocessors for cardiac arrythmia detection via peak detection on a single ventricular lead. Even when atrial leads are used, Intra Cardiac Electrogram (ICEG) morphology classification is required to separate some potentially fatal rhythms from harmless ones. [Leong and Jabri, 1992] The requirements of such a morphology classifier are:

- Adaptable to differing morphology within and across patients.

- Very low power consumption. ie. minimum energy used per classification.

- Small size and high reliability.

This paper demonstrates how this morphology classification may be done using a neural network architecture and thereby meet the constraints of the implantable arrythmia classification system. In addition, in loop training results will also be given for parity 4, another difficult nonlinear training problem.
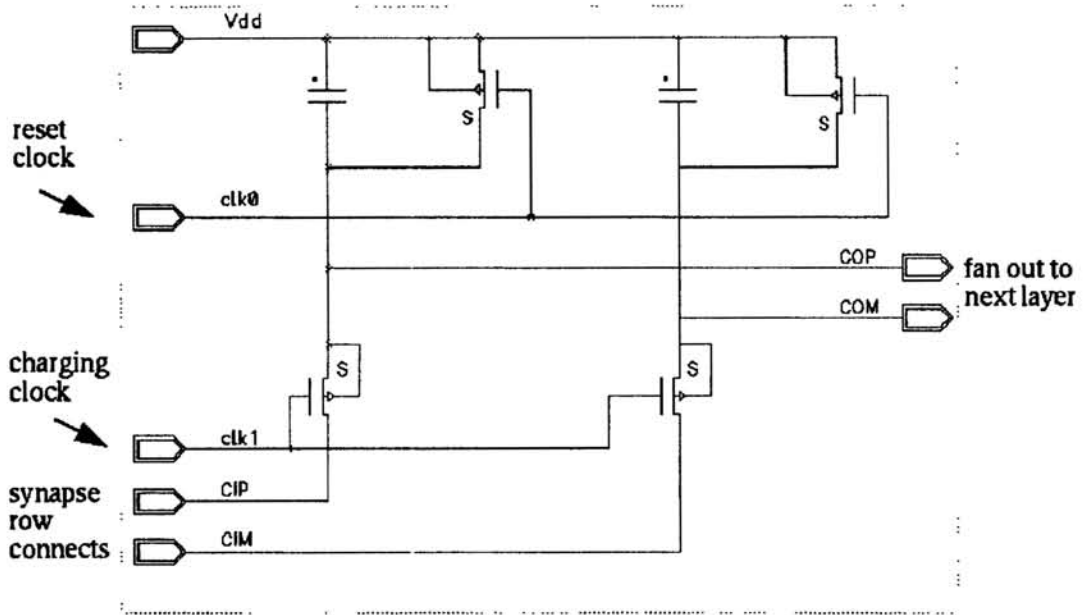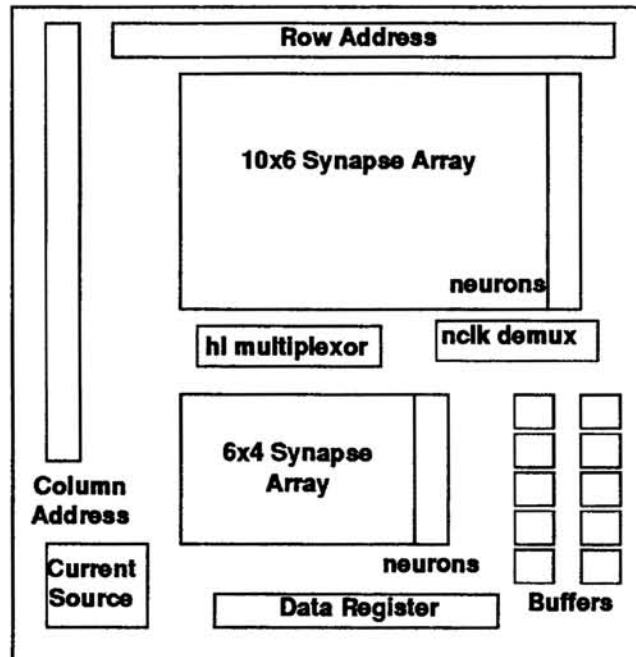
Figure 2: Wattle Neuron Circuit Diagram
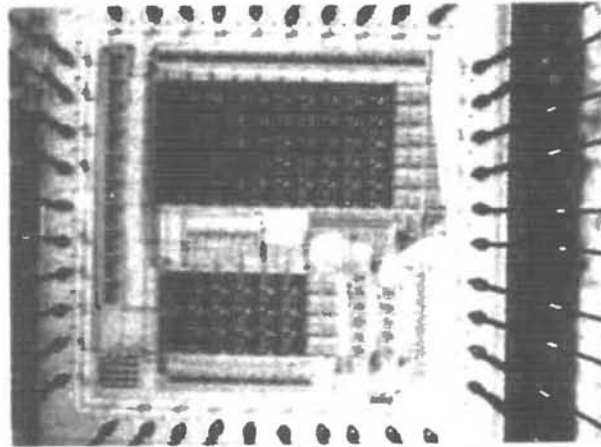

Figure 3: Wattle Floor Plan

Figure 4: Photomicrograph of Wattle

2   ARCHITECTURE

Switched capacitors were chosen for the neurons on Wattle after a test chip was fabricated to evaluate three neuron designs. [Coggins and Jabri, 1993] The switched capacitor design was chosen as it allowed flexible gain control of each neuron, investigation of gain optimisation during limited precision in loop training and the realisation of very high effective resistances. The wide gain range of the switched capacitor neurons and the fact that they are implemented on chip has allowed Wattle to operate over a very wide range of bias currents from 1pA LSB DAC current to 10nA LSB DAC current.

Signalling on Wattle is fully differential to reduce the effect of common mode noise. The synapse is a multiplying digital to analogue convertor with six bit weights. The synapse is shown in figure 1. This is identical to the synapse used on the Kakadu chip [Leong and Jabri, 1993]. The MDAC synapses use a weighted current source to generate the current references for the weights. The neuron circuit is shown in figure 2. The neuron requires reset and charging clocks. The period of the charging clock determines the gain. Buffers are used to drive the neuron outputs off chip to avoid the effects of stray pad capacitances.

Figure 3 shows a floor plan of the wattle chip. The address and data for the weights access is serial and is implemented by the shift registers on the boundary of the chip. The hidden layer multiplexor allows access to the hidden layer neuron outputs. The neuron demultiplexor switches the neuron clocks between the hidden and output layers. Figure 4 shows a photomicrograph of the wattle die.

3   ELECTRICAL CHARACTERISTICS

Tests have been performed to verify the operation of the weighted current source for the MDAC synapse arrays, the synapses, the neurons and the buffers driving the neuron voltages off chip. The influences of noise, offsets, crosstalk and bandwidth of these different elements have been measured. In particular, the system level noise measurement showed that the signal to noise ratio was 40dB. A summary of the electrical characteristics appears in table 1.

Table 1: Electrical Characteristics and Specifications

Parameter	Value	Comment
Area	$2.2 \times 2.2 mm^2$	
Technology	1.2um Nwell CMOS 2M2P	standard process
Resolution	weights 6bit, gains 7bit	weights on chip, gains off
Energy per connection	43pJ	all weights maximum
LSB DAC current	200pA	typical
Feedforward delay	1.5ms	@200pA, 3V supply
Synapse Offset	5mV	typical maximum
Gain cross talk delta	20%	maximum

A gain cross talk effect between the neurons was discovered during the electrical testing. The mechanism for this cross talk was found to be transients induced on the current source reference lines going to all the synapses as individual neuron gains timed out. The worst case cross talk coupled to a hidden layer neuron was found to be a 20% deviation from the singularly activated value. However, the training results of the chip do not appear to suffer significantly from this effect.

A related effect is the length of time for the precharging of the current summation lines feeding each neuron due to the same transients being coupled onto the current source when each neuron is active. The implication of this is an increase in energy per classification for the network due to the transient decay time. However, one of the current reference lines was available on an outside pin, so the operation of the network free of these transients could also be measured. For this design, including the transient conditions, an energy per connection of 43pJ can be achieved. This may be reduced to 10pJ by modifying the current source to reduce transients and neglecting the energy of the buffers. This is to be compared with typical digital 10nJ per connection and analogue of 60pJ per connection appearing in the literature. [Delcorso et. al., 1993], Table 1.

4   TRAINING BOTH GAINS AND WEIGHTS

A diagram of the system used to train the chip is shown in figure 5. The training software is part of a package called MUME [Jabri et. al., 1992], which is a multi module neural network simulation environment. Wattle is interfaced to the work station by Jiggle, a general purpose analogue and digital chip tester developed by SEDAL. Wattle, along with gain counter circuitry, is mounted on a separate daughter board which plugs into Jiggle. This provides a software configurable testing environment for Wattle. In loop training then proceeds via a hardware specific module in MUME which writes the weights and reads back the analogue output of the chip. Wattle can then be trained by a wide variety of algorithms available in MUME.

Wattle has been trained in loop using a variation on the Combined Search Algorithm (CSA) for limited precision training. [Xie and Jabri, 1992] (Combination of weight perturbation and axial random search). The variation consists of training the gains