FOSTER STOCKBROKING

EQUITY RESEARCH

Information Technology

19 May 2016

BrainChip Holdings Ltd (BRN)

Huge potential for neural computing in the age of artificial intelligence

Event:

- We initiate research coverage of BrainChip Holdings Ltd (BRN).

Investment Highlights:

- BrainChip (BRN) has developed a chip-embeddable technology which has the ability to learn autonomously, i.e. without any prior software instructions. This so-called Spiking Neuron Adaptive Processor (SNAP) is essentially a digital version of the way the brain operates, and consists of digital circuits that simulate the function of biological neural cells, including dynamic synapses, dendrites and somas. It is able to process hundred million input events per second, which it evaluates autonomously enabling it to learn and classify patterns by itself.
- Similar to the human brain, this Neural Computing technology enables incoming information to be processed in parallel, i.e. all neurons working simultaneously and associatively. In contrast, today's mainstream computer chips operate sequentially, i.e. the chip has to perform and complete a certain calculation according to specific software instructions before it can move onto the next calculation. Consequently, the SNAP technology is many times faster than the human brain and supercomputers, such as IBM's Blue Gene/Q, which comprises tens of thousands of individual processors. Furthermore, SNAP is substantially more energy efficient making it highly suited to mobile battery operated devices.
- SNAP is an embeddable technology, which means that it can be integrated into customers' specific chip designs. SNAP learns autonomously, but in many cases the technology needs to be configured for a specific task. At this time BRN is offering design services to help customers design the technology into their products. In the near future BRN will provide customers with a development kit that enables them to customize SNAP for a specific application, such as voice recognition on mobile phones, further increasing efficiency of the chip.
- Rather than manufacture the chips itself, which is highly capital intensive, BRN aims to commercialize its technology through licensing agreements with chip manufacturers, which typically include up front, non-recurring engineering and licensing fees and recurring royalties per manufactured chip once it goes into production.
- Apart from the speed advantage, the fact that a SNAP can recognize and classify real world inputs autonomously is a major breakthrough in the world of computer processing and Artificial Intelligence in particular, which is a big advantage over 'Deep Learning' systems that need weeks to train. The technology can be applied to many fields, including autonomous vehicles, finance, security, connected devices in the Internet of Things and industrial applications.
- The individual segments of BRN's addressable market, including autonomous robots, smart vision systems, smart machines, driverless cars and the relevant sub segments of the Internet of Things, are all expected to show exponential growth in the next five to ten years to reach tens of billions of dollars in market size.

Recommendation:

- We start our research coverage of BRN with a Speculative Buy recommendation, based on the company achieving milestones that de-risk the commercialisation of its technology. We expect an announcement around an initial licensing deal sometime in 2016, which should be a major share price catalyst.

Rating	**Speculative Buy**
Previous	n/a
Price Target	**n/a**
Previous	n/a
Risk	Very High

Share price	**$0.140**
ASX code	BRN
52 Week Low-High	$0.115-0.459

Pro-forma Capital Structure*	
Ordinary shares (M)	733.5
Market Cap (A$M)	**103**
Net Cash/(Debt) (A$M)	5.2
EV (A$M)	**97**
Performance shares	87.0
Options	29.6
Diluted mkt cap (A$M)	**119.0**
12mth Av Daily Volume ('000)	3,407
Assumes rights issue is fully placed	

Board of Directors	
Mick Bolto	Non-Executive Chairman
Peter Van Der Made	Chief Executive Officer & CTO
Neil Rinaldi	Non-Executive Director
Adam Osseiran	Non-Executive Director

Major Shareholders	
Peter van der Made	18.9%
Robert Mitro	13.9%

Share Price Graph A$

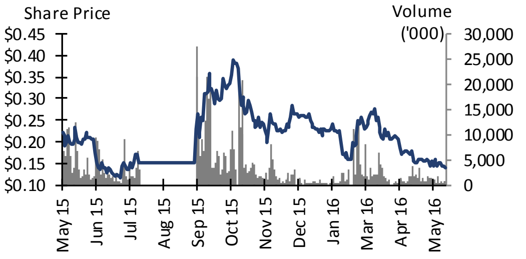

Foster Stockbroking is engaged by BRN to provide corporate services for which it has earned fees and continues to do. Services include, but are not limited to, research, marketing roadshows, and capital markets advisory. Foster Stockbroking is advising on BRN's 1-for-26 rights issue at $0.15 announced April 2016 for which it is earning fees.

ARTIFICIAL INTELLIGENCE … NOT THAT INTELLIGENT, REALLY

Artificial Intelligence is pre-programmed intelligent behavior

- Artificial Intelligence (AI) is the science and engineering of making intelligent machines, especially intelligent computer programs, where intelligence is the computational part of the ability to interpret the world, and perform actions in the world. In other words, AI relates to a pre-defined set of instructions a computer program must follow to achieve an optimal outcome. For instance, a Google driverless car must be programmed to slow down if its sensors say it is getting too close to the car in front. While this behavior may seem intelligent, it is actually "just" pre-programmed.

- Consequently, AI is not to be confused with computers or machines having an understanding of what they are actually working on, or a mind or even a soul, as depicted in some science fiction movies. It simply refers to pre-programmed intelligent behavior.

Artificial Intelligence today depends on massive computer power

AI today is not yet practically applicable to everyday problems

- The level of intelligence in an AI machine depends on the ability of software designers to capture all or most situations this machine may encounter as well as the speed at which the machine is able to process the incoming information and respond to it.

- Today's state-of-the-art AI projects typically involve massive computer power in an effort to approach the speed at which the human brain would solve the same problem. For instance, an IBM Blue Gene/Q supercomputer has been used in AI research. Depending on the configuration, a Blue Gene/Q can comprise more than 1.6M individual processor cores, filling entire floors (see Figure 1). This machine was at least 5,000x *slower* than the human brain: 1/3 the size of the brain was emulated in a simplified model, which was 1,542x slower than real-time (Source: IBM).

- Obviously, this is not a workable solution for many AI applications, e.g. to help a drone avoid moving obstacles or to move a self-driving car through traffic.

Figure 1: IBM's Blue Gene/Q

Source: IBM, Foster Stockbroking.

LEARNING VERSUS PROGRAMMED, RULES-BASED BEHAVIOUR

Actual learning allows computers to respond to an infinite number of circumstances

- When it comes to AI, a clear distinction should be made between actual learning and rules-based behavior. In the earlier example of the driverless car, rules-based programming dictates how the computer treats incoming information and how it responds to it.

- However, ideally we would want a computer to be able to learn from various inputs and its own subsequent responses to these inputs so that it doesn't make the same mistake twice. For instance, if the programming in a driverless car is not adjusted after it crashes into a traffic light during a test run, it will continue to make the same mistake in subsequent test runs.

- Computers that learn from their mistakes will not be bound by the finite number of circumstances that their programmers have been able to capture in the software that runs the computers. Instead, computers with the ability to learn will be able to adapt to an infinite number of circumstances that the software developers may never have thought of in the first place.

- One approach to self-learning computers has been to develop a processing architecture that mimics the spiking neural network inside the human brain.

ARTIFICIAL NEURAL NETWORKS ARE A STEP UP FROM RULES-BASED PROGRAMMING

- Artificial Neural Networks (ANNs) are computational models inspired by the structure and functions found in biological neural networks, e.g. human and animal brains. ANNs consist of interconnected neurons that exchange information, starting with the input coming in from the outside, e.g. sound or images.

ANNs can learn by tuning the weight of a signal

- The connections between the individual neurons have numerical weights that determine the strength, or importance, of the signal between the neurons. These weights can be tuned based on prior experience enabling the ANN to learn by adjusting the weights over a period of time.

- The ability of ANNs to learn, e.g. visual or acoustic patterns, based on inputs rather than through extensive computer programs, is what differentiates ANNs from rules-based programming AI.

- However, ANNs also come in various forms. Individual neurons in software-based ANNs still need to be tuned by humans to achieve the best results, whereas a hardware-only ANN, such as the one BRN has developed, can learn **fully autonomously**, which is a very important distinction.

Figure 2: Artificial Intelligence and Neural Networks

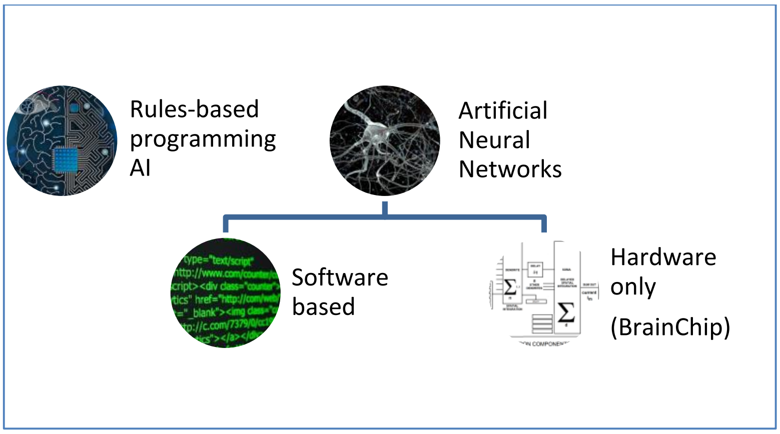

Source: Study.com, dglopa.com, Foster Stockbroking estimates.

NEURAL NETWORKS 101

To get a better understanding of ANNs, we first need to understand what it is they are actually trying to mimic, i.e. the biological neural network in the human brain.

The working of neurons in the human brain

86BN neurons in the human brain

- The average human brain consists of approximately 88 billion neurons, or nerve cells, and 100 trillion synapses that contain multi-level memory. Neurons are cells that process and transmit information through electrical and chemical signals (see Figure 3), also known as spikes. Sensory neurons respond to outside signals, such as sight, touch and taste and pass on these spikes through their axons and synapses to adjacent neurons.

Figure 3: An actual and stylized biological neuron

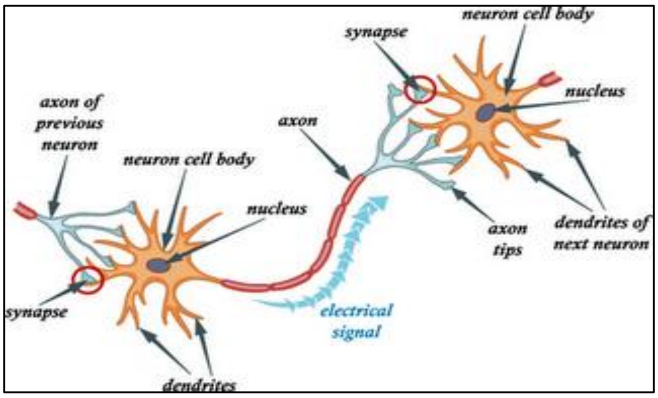

Source: Study.com, dglopa.com, Foster Stockbroking estimates.

Synapses are pathways between neurons

- **Synapses** act as *pathways* between neurons and will fire off a chemical output signal (neurotransmitters) to the next neuron if the aggregated signals coming in through its dendrites, i.e. from other neurons, exceed a certain voltage threshold (voltage gradient). In other words, if the incoming signals combined are strong or relevant enough to pass on. Synapses have memory and remember events that occurred previously. The synapse strength is modified as a result of learning.

  Dendrites receive nerve signals

  Synapses come in two main forms, **inhibitory synapses** and **excitatory synapses**. The former pass on an inhibitory signal. The latter increase the likelihood that the next neuron fires off a spike. In turn, the adjacent neuron will receive these incoming nerve signals through its synapses, connected to **dendrites**, which extend from the cell body and can branch out multiple times effectively forming what is called a dendritic tree. Dendrites receive and transmit the electrochemical stimulation received from other neurons to the cell body, or soma.

The axon hillock plays a major role in actual learning

- The **axon** extends from the soma and carries electrical nerve signals and neurotransmitters to the neuron's outgoing synapses, where the process repeats in the next neuron. Axons can grow as long as 1 meter in the human body and emerge from an area called the **axon hillock**, which plays a vital part in the actual functioning of neurons, i.e. where the neuron decides whether incoming spikes from another neuron are strong enough, as measured in a voltage, and should be passed on to the next neuron.

We will discuss the latter in more detail below as this feature goes to the heart of learning, both in biological neurons and ANNs.

- The **soma**, or cell body, is essentially the engine room of the neuron, producing a number of different proteins for the dendrites, axons and synapses as well as energy for the cell to function.

- Neurons come in three basic varieties, including sensory neurons as discussed above. Additionally, based on signals from the brain and spinal cord, motor neurons send signals to muscles to generate a muscle contraction, and thus movement. Lastly, interneurons simply connect neurons to other neurons in the same region of the brain or spinal cord. Each neuron performs a computation, integrating the strengths stored in the synapses and the activity of input spikes.

The neural workflow

"Action Potential" triggers chemical signal in the synapses

- In the neural "workflow", a neuron receives electrical signals, or spikes, through its synapses and dendrites. Based on the aggregated signal strength and whether this electrical signal reaches a critical voltage threshold, the axon hillock computes whether or not to fire off a spike of its own through its axon towards the synapses. If it does, this spike, called **Action Potential,** triggers a chemical signal in the outgoing synapses (Figure 4) that ejects neurotransmitters towards the next neuron where the process repeats itself. Neurotransmitters are converted to electrical signals of different potential levels and persistence in synapses.

Figure 4: Transfer of information through synapse

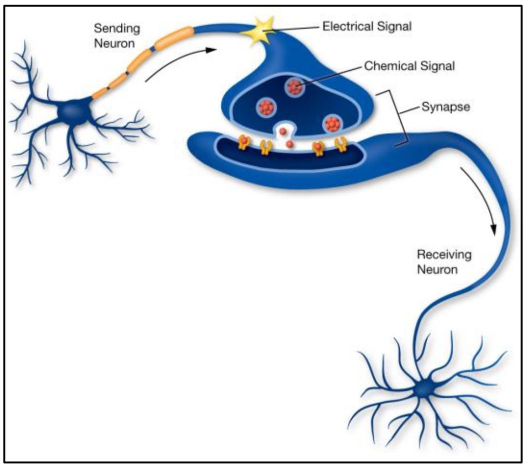

Source: Foster Stockbroking estimates, learn.genetics.utah.edu

THE BRAIN REGARDS ALL SIGNALS AS PATTERNS

- It is important to note that neurons regard all incoming signals as patterns, regardless of the nature of the signal, e.g. signals triggered by sight, touch, taste, sound etc. are all converted into spikes before entering the nervous system.

The brain learns through pattern recognition

- This is important to understand as it implies that the human brain essentially reduces all analogue information into a universal signal. Based on the coincidence of spike patterns, the brain works out what information is relevant and what is not, i.e. it learns based on pattern recognition.

BRN's approach to chip design is based on this principle. However, before discussing BRN's approach to chip design it is important to understand the mechanism by which ANNs currently learn, i.e. what happens inside the artificial neuron and which methods are used to get ANNs to learn, i.e. supervised versus unsupervised learning.

1. THE LEARNING MECHANISM IN ARTIFICIAL NEURAL NETWORKS

- In ANNs incoming spikes (x_1, x_2, ….) from other artificial neurons are given a weight (w_1, w_2, ….) as a way to attribute importance to individual spikes (synaptic weights) (Figure 5). The weighted spikes are subsequently summed and result in a total input signal coming in to this particular neuron's core. The aggregated spike is transferred to the neuron's **Activation Function**, which subsequently determines if the sum of all signals is relevant enough and needs to be passed on to other neurons in the network.

- There are various activation functions, for instance a threshold activation function: if the sum of all the incoming spikes exceeds a certain value, the neuron fires off a spike of its own. If the sum is lower than the threshold, no spike will be fired by the neuron.

Figure 5: Schematic of an Artificial Neural Network (ANN)

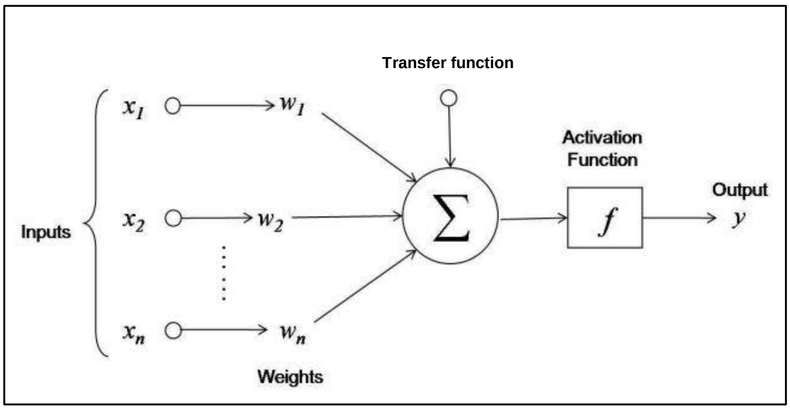

Source: stackexchange.com, Foster Stockbroking estimates.

Learning by adapting the weights

- Some spikes are more relevant than others in recognizing patterns so the input weights are adaptable based on the feedback the neuron receives to its own output spikes. Consequently, artificial neurons learn by adjusting the weights attributed to each input spike and the feedback they get on their subsequent output spikes.

- For instance, if the input spikes from a particular neuron or group of neurons does not seem to lead to improved results, the weights attributed to signals coming from those neurons are lowered in a trial-and-error method, similar to how a biological brain works.

The weights effectively form the memory of the neuron

- Consequently, the weights in each synapse says something about the past, i.e. it effectively forms the memory of the neuron, similar to how biological neurons have a memory in the form of the voltage gradient discussed earlier.

2. SUPERVISED VERSUS UNSUPERVISED LEARNING

In order to train ANNs, sets of **training data** are presented to the network.

Working towards a known outcome

- In supervised learning, both the input and required output data are known by the ANN and the learning process is based on comparison between current output and expected output. The weights of the input spikes are corrected over a period of time in a trial-and-error method with the objective to gradually minimize the variance between the actual output and expected output.

Adjusting the weights through trial-and-error

- For instance, an ANN may be trained for facial recognition and is fed a set of input spikes and the picture of a person (the output) it needs to derive from this input. Through trial-and-error the individual neurons in the network will continue to adjust the input weights until the network output resembles the picture.

- After this process, the neurons in this artificial network will have been optimized for facial recognition and it should take the ANN substantially less time to recognize a second face when presented with a new set of input data.

In the next step, the ANN is presented with **validation data**.

- To verify the new model, the weights are checked for accuracy in the model prediction. No adjustment are made to the weights but based on the results, the training procedure and parameters may be further optimized, for instance to achieve higher precision.

Finally, after training and validation are completed, the ANN is given **test data.**

- The model will be tested with new, independent, input data to check the performance. Accuracy of the ANN will depend on the magnitude of the error between the desired output (e.g. a face in the facial recognition example above) and the ANN's actual output.

The process above describes what is known as **supervised learning**, i.e. the end result is known and the ANN needs to tweak its neurons in a repeated process to the point that the ANN's output closely resembles the desired output that was known beforehand, i.e. the error is minimized. This is how most ANNs work today.

The neural network needs to learn on its own

- In **unsupervised learning** no expected output pattern is given. The ANN is only given a set of input data and needs to start recognizing and classifying the data patterns by itself. Through the learning process, the weights are adjusted by an autonomous function in the neuron that is independent of the target output. The ANN learns from its own experiences.

TODAY'S ANNs ALL HAVE SOFTWARE TELLING THEM HOW TO LEARN

By default software limits the ANN's learning capabilities

- While the learning method is different for supervised and unsupervised learning, today's ANNs will need a learning algorithm no matter which method is chosen, i.e. a software program that tells the ANN how it needs to adjust the synaptic weights.

- There are multiple learning algorithms available, description of which goes beyond the scope of this report. Suffice it to say that any learning algorithm is programmed by a software developer, which inherently limits and correlates the learning capacity of the algorithm to the intellectual capacity of the developer and the programmed scope of the software.

- Put differently, if a network were to learn independently, i.e. without parameters set by humans, not only might the learning process take place substantially faster, the system might also "connect dots" that humans wouldn't have programmed the network to find. Or connect dots in a totally different fashion than a human programmer would.

In other words, software-based autonomous learning is inherently limiting the possible outcomes. A hardware-only solution could potentially be substantially more efficient, both in time and potential applications. Enter BRN.

BRN'S SOLUTION IS HARDWARE-ONLY

- BRN's approach to ANNs differs from today's existing ANNs in several very significant ways. Most importantly, BRN's solution is 100% hardware-based, i.e. there is no pre-programmed function that tells the artificial neurons how to adjust the synaptic weights and there is no activation function that tells the neurons when to fire off a spike.

Milestone 3 demonstrated autonomous learning

- In other words, in the absence of algorithms programmed by humans, all the learning in a BrainChip is purely done by the ANN itself. This was demonstrated recently when BRN achieved Milestone 3, which involved a BrainChip autonomously processing and classifying visual images from a freeway overpass **(Figure 6**). The chip had no prior knowledge of what information it was receiving, but within seconds it was able to discern cars driving in separate lanes and it started keeping count of traffic numbers.

Figure 6: BRN's chip uses temporal signal inputs

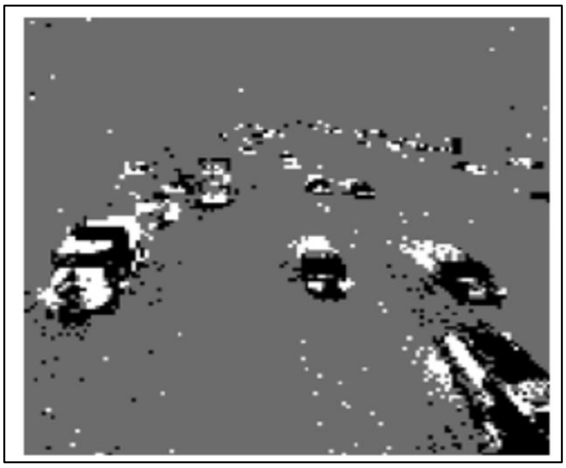

Source: Company, Foster Stockbroking estimates.

Other key differences with existing Artificial Neural Networks (ANNs)

- As opposed to software-based ANNs that can process one instruction at a time, spiking neurons in BRN's hardware-only ANN do not need to work their way through a full computational cycle to come to an outcome on which to base their next calculation. Because they process spikes (temporal signals, Figure 7) without having to calculate anything, BRN's ANNs have a tremendous speed advantage.

BRN's chip is substantially faster and more energy efficient

- Further contributing to its speed advantage is BRN's parallel architecture, i.e. information is processed by all neurons at the same time. This enables BRN's ANN to perform massive parallel processing, whereas today's ANN's process information sequentially. The learning process in a BrainChip takes place within seconds, whereas today's ANN's may take up to 50x to 80x longer to achieve the same result.

- Because of their size requirements, today's ANNs consume substantially more energy than BRN's solution, making them unsuitable for most mobile applications. BrainChip, however, is ideally suited for mobile applications, such as speech, face and handwriting recognition in mobile phones as the energy consumption is just 1,000th of that of a traditional GPU or CPU.

Figure 7: BRN's chip uses temporal signal inputs

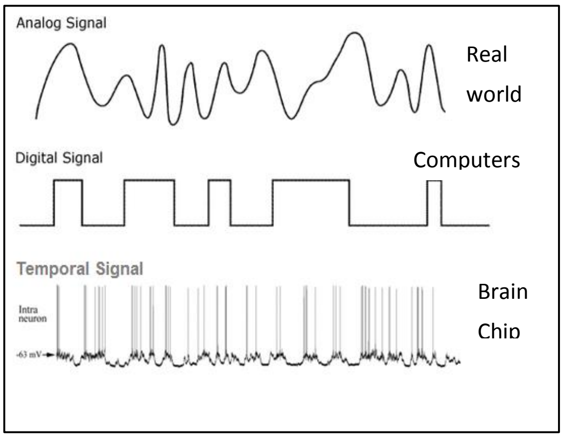

Source: Company, Foster Stockbroking estimates.

THE HEART OF A BRAINCHIP: SPIKING NEURON ADAPTIVE PROCESSOR (SNAP)

- BRN's Spiking Neuron Adaptive Processor (SNAP) can comprise up to ten thousand individual artificial neurons that process input spikes directly in hardware and are all updated in parallel. Sensory neurons convert physical stimuli, e.g. images, sounds etc into spikes. As described earlier, learning occurs based on input intensity, i.e. many spikes in a short period of time, or repetition through feedback loops at each neuron.

BRN's artificial neurons also learn from their own feedback

- As can be seen in Figure 8 a BrainChip closely resembles the architecture and process of a biological neuron. Spikes coming into the synapses (left hand side) of a particular dendrite are aggregated in the dendrite and subsequently forwarded to be aggregated at the neuron level, i.e. signals coming in from the artificial neuron's multiple dendrites are aggregated. The summed signal is passed on to the axon hillock where the output spike (the Action Potential) is generated.

- The output spike is sent to all connected neurons as well as being fed back into the neuron's own synapses. This so-called self-feedback loop attempts to incorporate time and memory into the network, i.e. the neuron learns from its own output, similar to the biological brain. These ANNs are also known as recurrent networks and have a higher ability to learn compared with feedforward networks in which the signals only move in one direction, i.e. only to other neurons without any feedback into the neuron itself.

Figure 8: BrainChip's SNAP mimics biological neurons

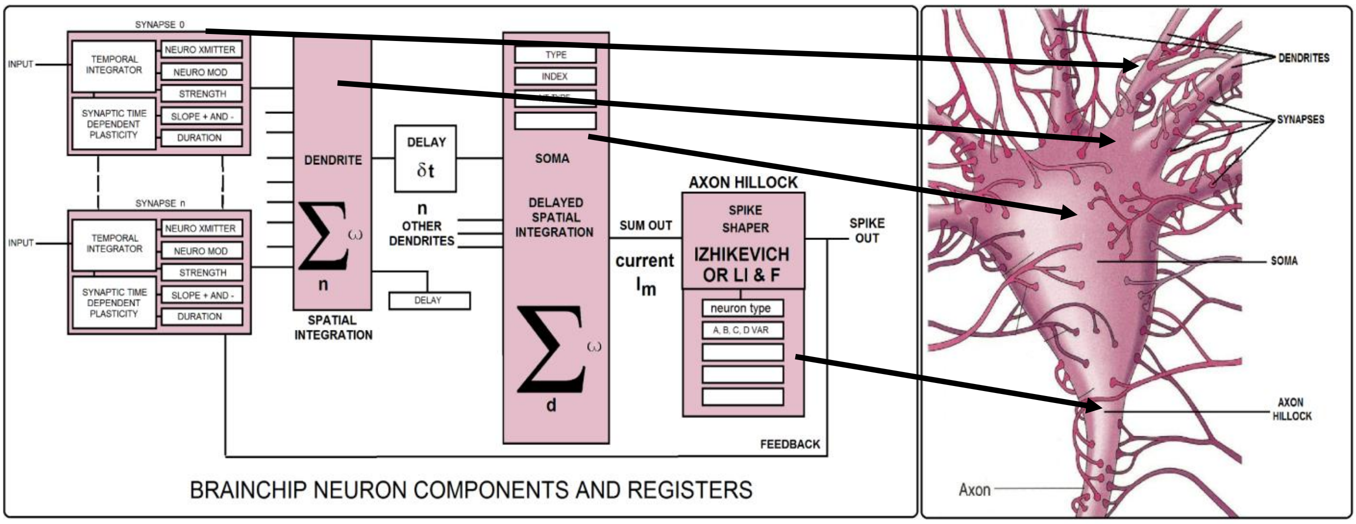

Source: Company, Foster Stockbroking estimates.

AUTONOMOUS LEARNING THROUGH SPIKE TIMING DEPENDENT PLASTICITY

- Spike Timing Dependent Plasticity (STDP) is a process that adjusts the strength of connections between neurons in the brain, as we described previously. The nature of the adjustment, i.e. stronger or weaker, depends on the timing of input spikes coming into a particular neuron and its own output spikes in response to that input.

- If there is a short time lag between input and output spikes the connection between the two relevant neurons will strengthen or vice versa if the time lag is longer. In other words, the plasticity (rate of adjustment) of the connection depends on timing of the spikes.

The synaptic weights form the memory of the neuron

- These characteristics of an artificial neuron effectively form the memory of the neuron. The neuron knows it previous response to a signal from another neuron based on the weight their connection has and uses this information in its future responses.

The combination of STDP and BRN's hardware-only architecture, i.e. without any human programming involved, is what facilitates the autonomous learning process inside a BrainChip.

COMPARABLE SOLUTIONS STILL REQUIRE HIGH DEGREE OF HUMAN INVOLVEMENT

Several other neural networks that have been developed in recent years still require either software code or human interaction to learn, i.e. learning is still not autonomous in these chips.

No autonomous learning capability in alternative technologies

- IBM's True North chip is a neural network with more than 1M neurons on board. The 256 synapses on each neuron are programmable, i.e. need to be told what to do in varying circumstances by a human programmer. As explained previously, this implies that the chip essentially will only be able to handle situations that the software programmers anticipated, limiting the learning capability of the chip in real-life.

- Qualcomm's Zeroth neural network is designed to run on a future version of the company's Snapdragon system-on-a-chip (SOC) currently in operation on millions of mobile phones. However, most of the learning will need to take place "off line" on large deep learning computer networks after which time the "learnings" are transplanted or programmed into individual chips. Additional learning once the chip is being used by consumers will be very limited compared to true autonomously learning chips.

Deep learning requires huge processing capacity

- Cognimem have artificial neural network products on the market comprising of either 1,024 or 2,048 neurons, which when stacked can reach a maximum capacity of 8,192 neurons. These chips are trainable neural networks, but require training for a very specific purpose prior to being deployed.

- Deep learning projects run by giants such as Google and Facebook, require large scale computer power to run and learn, making them impractical for many applications, such as mobile devices, at least for now. Additionally, they run on algorithms rather than being hardware-only. And while these algorithms may be very advanced, i.e. Google's DeepMind beat one of the world's best Go players recently, to a certain extent they are still bound by the scope of the program.

WIDE RANGE OF APPLICATION AREAS

The application areas for BRN's technology are plentiful, in our view, in particular in areas where activities are non-repetitive and require on-the-fly adjustments and responses, which to date only humans can make efficiently:

- **Autonomous vehicles/driverless cars** that need to adapt to continuously changing circumstances, sometimes within split seconds.
- **Robotics**, where autonomous learning will help robots interact with the outside world.
- **Exploration and mining**, which both require responses to local circumstances in the absence of humans, e.g. underground, deep sea and space exploration.
- Forecasting, such as **weather forecasting** and **financial modelling** that involves processing of very complex input data and patterns that take current ANN's hours or even days to process.
- **Security monitoring** and **control systems**, where pattern recognition plays a dominant role within the overall system, i.e. autonomous recognition of defects before they occur.
- **Sniffing** and **listening devices**, e.g. for border control and fault detection, require very granular approach to data analysis and pattern recognition. For instance, very small changes to the pitch of a stationary airplane engine may hint at looming defects.
- **Speech** and **Image Recognition** are very obvious application areas, especially in mobile devices, conditional access and voice controlled systems.
- **Eye and Ear implants** process real world information and would benefit tremendously from more autonomous learning and pattern recognition to better cope with different real world circumstances.
- The fields of **Biotechnology** and **Genetics** require advanced pattern recognition and associative computing capacity for R&D purposes.
- Currently about 60% of computer processing time in **Computer Gaming** consoles, like the XBOX and PlayStation, is used up by genetic algorithms and neural networking to control things like computerized opponents' AI responses to the human players' actions. Reducing the number of cycles and accelerating cycle time of the CPU by introducing a hardware-only neural networks will dramatically increase game console speeds, which we believe will be a very valuable application area for BRN.

ADDRESSABLE MARKET GROWING EXPONENTIALLY

Neuromorphic chip market to reach US$4.8B in 2022

- There are various definitions of BRN's addressable market that overlap or complement each other. The narrowly defined market for neuromorphic chips, i.e. semiconductors like the BrainChip and software-based ANNs, is estimated to reach US$4.8B by 2022 according to MarketsAndMarkets compared to approximately US$500M in 2014.

Broader addressable market

- However, we believe this severely understates the market potential for neuromorphic-like functionality, such as BRN provides, because BRN's technology is likely to be integrated or embedded into customers' existing architecture as well as be sold stand alone. For instance, mobile handset manufacturers might want to include a limited version of a BrainChip on their mobile phones, just to unlock the phone through speech recognition. This would require just a limited number of artificial neurons to be integrated in the chipset, rather than a full-fledged neuromorphic chip on board. We will elaborate on this below.

Suffice it to say that we need to look at a wider definition of the addressable market for autonomous learning than just neuromorphic chips to get a sense of the market potential for BRN.

- The market for Smart Machines in general totaled US$6.3B in 2014 and is expected to exceed US$15B in 2019. By 2024 this market is expected to be worth US$41B[1], implying a CAGR of 21% over a ten year period (Figure 9).

Figure 9: Neuromorphic chip market by application (in US$ M)

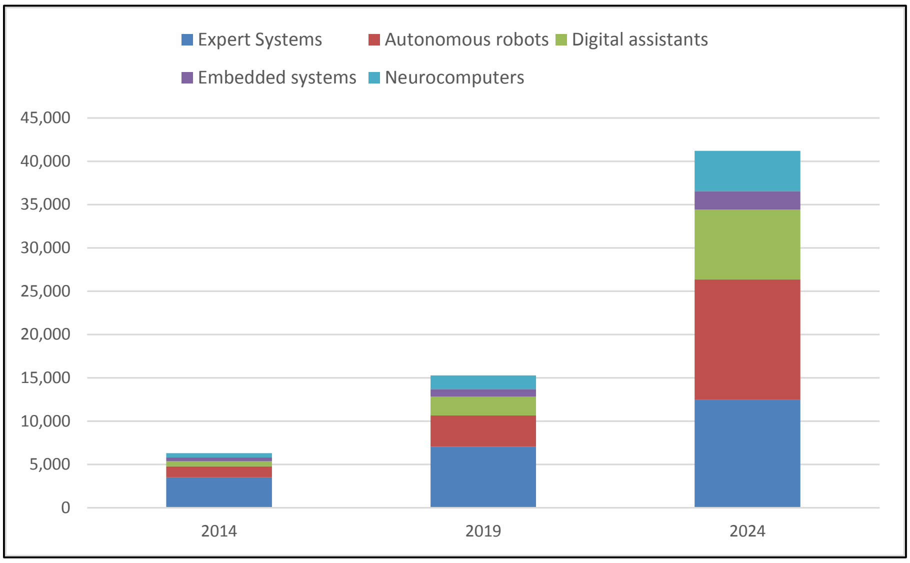

Source: BCC Research, Foster Stockbroking estimates.

- Within this market, the size of the autonomous robots segment was estimated at US$1.3B in 2014. It is expected to grow to US$3.6B in 2019 and US$14B by 2024, implying a 27% CAGR over a ten year period.

- Expert Systems[2] is another large component within the Smart Machines market, totaling US$3.5B in 2014, growing to US$12.4B by 2024, or 13.5% CAGR over ten years.

Other application areas to show strong growth as well

Markets for robotics, autonomous vehicles and expert systems are large

- Boston Consulting Group estimates the market for autonomous vehicles at US$42B by 2025, with the average additional cost of the technology to facilitate autonomous driving estimated at US$2,000 to US$10,000 per car. These costs include the additionally required hardware, software, sensors etc. We would expect the neuromorphic chipset to account for less than US$100 per car.

- However, other functionality in the car, such as unlocking and starting the car by voice, facial or fingerprint recognition, will likely require neuromorphic chips as well. Additionally, this sort of functionality is already being included in cars sold today.

In other words, the addressable market for BRN in the automotive market already exists today.

- We believe this is also true for autonomous learning in mobile applications, for instance Siri in Apple phones as well as simpler, stand-alone functionality, such as unlocking phones through voice recognition.

- Another interesting market is Smart Vision. The global market for Smart Vision systems, such as production inspection camera's on conveyors, is expected to grow approximately 9% annually from US$8.1B in 2015 to US$12.5B by 2020, according to MarketsAndMarkets. While the US

[1] BCC Research, 2014

[2] Expert Systems are AI-based systems that convert the knowledge of experts in a specific subject into a software code. This code can be used to answer questions or to run queries submitted through a computer. E.g Woodside Petroleum has been using IBM's supercomputer Watson as an expert system by uploading many years of engineering data from testing, projects and employee communications and has subsequently been running queries to answer a range of complicated questions around technical issues.

currently accounts for half this market size, Asia-Pacific is expected to show the highest growth in this segment in the next five years.

THE INTERNET OF THINGS WILL BE A KEY OVERALL DRIVER FOR BRN

- While the exact neuromorphic chip content in connected devices by 2020 is hard to quantify in monetary terms, the sheer number of devices connected through the Internet of Things (IoT), or the Internet of Everything (IoE), does give a rough idea of the addressable market the IoT presents to BRN.

50 billion connected devices by 2020

- Cisco estimates approximately 50B devices will be connected to the internet by 2020. This includes a very wide range of applications, such as sensors and detection systems (on oil platforms, smart energy meters, in driverless cars, in mobile devices, in pollution and fire detectors etc), smart cities (connected parking garages, traffic management, smart roads and lighting etc), retail and supply chain applications (NFC payment, inventory tracking), logistics and industrial control etc. The list seems endless.

Autonomous data classification before transmission

- Many of these applications will require more functionality than just sensing and transmitting. For instance, a particular sensor may need to qualify the incoming signals first before taking action, rather than just transmitting every signal coming in, such as a sensor in a high pressure valve that can recognize unusual pressure fluctuations. This sort of functionality is where neuromorphic chips come into their own, especially ones that can truly learn autonomously.

BUSINESS MODEL SIMILAR TO THAT OF OTHER SEMICONDUCTOR IP PROVIDERS

Well-understood licensing model

- In terms of commercialization, BRN will be using a licensing model that will be very familiar for prospective customers as it has been used for many years by chip design companies such as ARM Holdings and a wide range of other semiconductor companies, including fabless manufacturers like Qualcomm and Broadcom.

Figure 10: BRN revenue model caters to multiple customer types

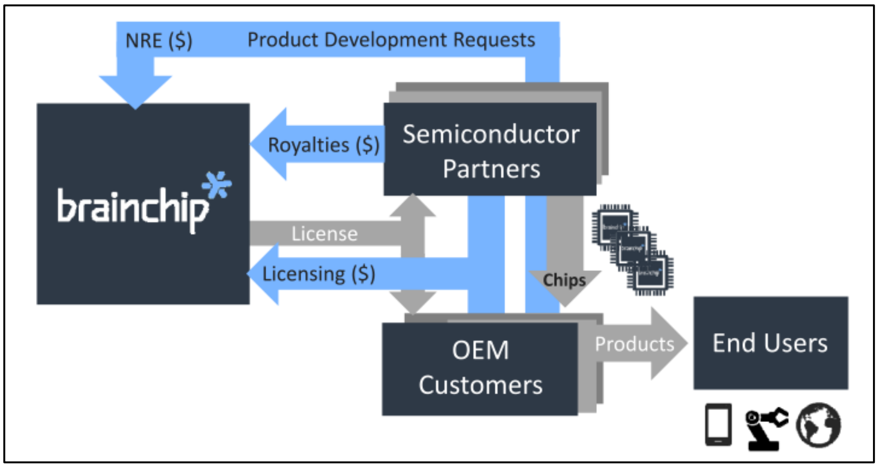

Source: Company, Foster Stockbroking estimates.

- Without manufacturing chips themselves, these companies design semiconductors and license out the Intellectual Property (IP) to customers, such as semiconductor manufacturers, mobile phone makers, consumer electronics companies etc. Revenues typically comprise of a one-off license fee and recurring royalty payments (see figure 10).

License fees and royalty income

- License fees can range anywhere from US$0.5M up to US$10M, depending on the application, anticipated production volumes, exclusivity in a specific market segment etc and are paid in several tranches. The first tranche usually amounts to 25% - 50% of the total license fee with the remainder paid in one or two tranches halfway and at the end of the engineering process.
- Royalty revenues are typically a percentage of the value of the sales price of the end product, usually somewhere between 2.5% and 20%, again dependent on the application, volumes etc.

Figure 11: Typical IP license revenue stream

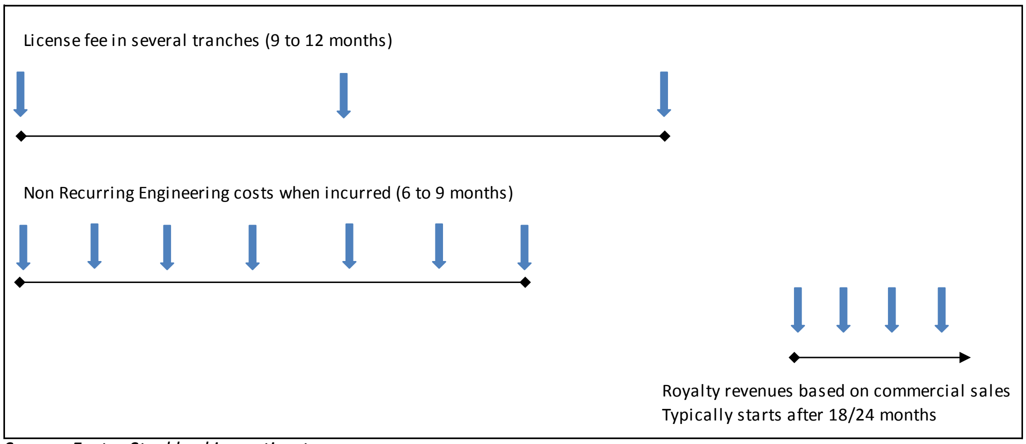

Source: Foster Stockbroking estimates.

Non Recurring Engineering costs to customize chip

- Furthermore, BRN will charge non-recurring engineering (NRE) costs to customers to adapt and customize the IP specifically for the customer's application of the IP. NRE can vary from a few hundred thousand dollars to several million, depending on how much of the application needs to be developed by BRN

Revenue stream lumpy initially

- Consequently, as BRN signs on customers we would expect to see a fairly lumpy revenue stream initially as NRE costs and license fees come in, followed by royalty income some 12 to 18 months later once the customer has fully integrated the BRN IP in its products and goes into commercial production.
- The timelines in figure 11 above provide a general indication of the time it will likely take from signing of an agreement to actual production. However, a lot will depend on the agility of the customer, i.e. how fast can they move, and the urgency to get the product to market. In other words, timelines could be shorter, e.g. 3 to 6 months for engineering and another 3 to 6 months to get into production.

BRAINCHIP WILL COMMERCIALISE THE IP IN DIFFERENT WAYS

We see several routes-to-market for BRN:

1) Integrate the IP into customers' existing IP library. E.g. chip design companies may want to include BRN IP into their various existing chip sets they will then sell to their own customers. This would likely be the fastest way for BRN to commercialise the technology as these companies are well established, trusted companies with large existing revenue streams.
2) Additionally, BRN will likely sell its IP directly to end customers, such as chip manufacturers, consumer electronics companies, mobile phone makers and OEM's that will need to do a bit of integration of BRN's IP into their existing chip architecture.
3) Revenue sharing partnerships with specialized semiconductor IP sales companies. These will accelerate the roll-out of BRN's IP.

THREE MILESTONES ACHIEVED

Milestones 1 to 3 mostly related to R&D and proof of concept

- To date BRN achieved three milestones that essentially evolved around R&D and proof of concept:
- **Milestone 1**: Simulating a race car demonstration in software for "proof of technology" by comparing BRN's SNAP technology to a traditional genetic learning algorithm technology. This demonstrated that the technology is practicable, is scalable and products can be verified using the tools that BrainChip develops.

- **Milestone 2**: Creation of a hardware-only version of its patented and proprietary autonomous learning Spiking Neural Adaptive Processor technology.

- **Milestone 3**: Developing and releasing a Client / Server Interface Tool to its autonomously learning SNAP technology enabling clients to develop their own solutions on top of the BrainChip hardware.

Milestone 4

Commercial milestone

- The fourth milestone is a commercial one: Execute an IP license agreement with an upfront payment of no less than US$500,000. BRN is in discussions with a number of prospects and expects to reach a commercial agreement with at least one sometime during 2016. It is unclear, however, what amount the company will receive as an upfront license payment.

IP PROTECTION THROUGH MULTIPLE GRANTED AND PENDING PATENTS

- In 2008 BRN was the first company to file a patent describing a learning digital neuromorphic chip design. To date it owns one granted patent[3] and has six filed and pending patents. The company's strategy is to file a number of additional patents going forward.

- However, while filing patents is a natural and required process for high-tech companies, it is not an ironclad defence, especially not for smaller companies. If one of BRN's patents (granted and filed) were to be challenged, it may prove very difficult for BRN to uphold the validity of the patent due to a potential lack of funds to properly defend it. The granted patent has been evaluated by a third party, and has proven to be a strong patent. The patent has been cited 25 times by companies such as IBM, QualComm, HRL and others.

ABSENCE OF ASX PEERS FOR VALUATION

- We have abstained from a peer group valuation for several reasons. Firstly, there are no comparable peers listed on the ASX. More importantly, the bulk of future revenues will comprise of royalty revenues. As these will take a few years to ramp up, there is no meaningful comparison base for a peer group valuation, which typically only looks out one or two years.

POTENTIAL FOR STRATEGIC TAKE-OVER IF TECHNOLOGY TAKE-UP IS STRONG

IP-related M&A activity may drive BRN share price after commercial roll out

- We believe there is clear potential for BRN to come into the sights of larger industry players if and when the take up of the technology is strong in the next few years. The semiconductor industry is prone to IP-related M&A activity. Examples include Broadcom being acquired by Avago for US$37B, implying a 4.4x P/Sales multiple and Intel buying Altera for US$16.7B.
- At the smaller end of town, companies like Synopsis (NASDAQ:SNPS) have been consistently snapping up semiconductor IP companies in recent years to expand their overall IP offering to large players in the semiconductor industry. Going forward, one of the key M&A drivers for semiconductor companies will remain to be the acquisition of technological breakthroughs, according to a recent KPMG survey, of which BRN might be one.

RECOMMENDATION: SPECULATIVE BUY

Initiate with Speculative Buy recommendation

- We start research coverage of BRN with a Speculative Buy recommendation.
- We believe BRN is will continue to derisk its technology by achieving commercial milestones over the next 12 months, and that the value of the business and equity will increase.
- Share price catalysts include signing of binding commercial agreements, higher than expected license fees and royalty percentages as well as faster than expected uptake by customers.

[3] Autonomous Learning Dynamic Artificial Neural Computing Device and Brain Inspired System, patent # 8,250,011.

SWOT ANALYSIS

Strengths

- The first hardware-only artificial neural network that is commercially available, bringing unique speed and cognitive advantages to market.
- Asset-light business model, well-understood by the industry.
- Strong technology endorsement from industry peers.
- Highly qualified management and advisory boards.

Weaknesses

- BRN has no existing commercial eco-system in place, unlike several potentially competing companies, such as Qualcomm and IBM. This may slow down initial commercial roll-out.
- Lack of revenues will require BRN to tap the market until such time as commercial license agreements are signed.

Opportunities

- Large and fast growing addressable market for neuromorphic chips.
- Potential to cooperate with established chip design companies.

Threats

- Rapid technological advancement by competing technologies/companies, bringing them at par with or leapfrogging BRN's technology.
- Patent infringement with potential lack of ability to enforce.

KEY RISKS

The following risks may negatively impact the BRN share price:

- **Funding risk.** In developing and marketing its products, BRN may need to raise further funds. There is a risk this could be dilutive to shareholders or may be difficult to achieve.
- **Technology risk**. The fast pace of technological development may result in BRN's technology becoming outdated or less efficient compared to competitors' solutions.
- **Key person risk.** BRN has a number of key people, such as Peter van der Made (acting CEO and CTO), that should BRN lose them, may weaken the company's management and strategy, adversely impacting earnings and future potential.
- **Earnings risk**. BRN is in the early growth stages of a company lifecycle and there is a risk that profitable earnings may not eventuate as soon as the market expects.
- **Reputation risk**. Any breach of security of BRN's online infrastructure may result in loss of credibility and reputation, potentially inhibiting the company's growth.

  Competitor risk. Competing companies providing similar products may erode BRN's market share, profitability and growth outlook.
- **Commercialisation risk.** Failure to commercially roll out the BrainChip technology and

Figure 12: Pro-forma Capital Structure (M)*

Shares on issue	M
Ordinary shares	733.5
Performance shares	87.0
Options	29.6
Fully Diluted	**850.1**

**Assumes rights issue announced April 2016 is fully placed.*

Source: Company; Foster Stockbroking estimates.

BOARD OF DIRECTORS

- **Mick Bolto**: **(Non-Executive Chairman and Director)** Mr. Bolto joined the Board in August 2015. He served as a partner at Mallesons for twenty years where he worked in mergers and acquisitions. He was instrumental in the structuring of and subsequent execution of numerous large-scale transactions in Asia, Australia, Europe and North America. Following his time at Mallesons, Mr. Bolto worked in private equity for a long period where he acquired extensive experience in creating strategy and business planning for SME's in order to ensure the delivery of viable business results.

- **Peter van der Made**: (**Executive Director, Chief Technical Officer and Acting Chief Executive Officer)** Mr. van der Made is the developer of BRN's core technology around artificial neural networks and joined the Board in September 2015. He has been at the forefront of computer innovation for 40 years. He is the inventor of a computer immune system at vCIS Technology where he served as CTO, and then Chief Scientist when it was acquired by Internet Security Systems, and subsequently IBM. Previously, he designed a high resolution, high speed colour graphics Anatomy chip for IBM PC graphics. Most recently he published a book, Higher Intelligence, which describes the architecture of the brain from a computer science perspective.

- **Neil Rinaldi: (Non-Executive Director**) Mr. Rinaldi joined the Board in May 2013. Based in Perth, he spent the previous 10 years in London culminating in his position as Managing Director of Truestone Capital, a London based corporate advisory group that specializes in the provision of corporate advice ranging from M&A activity through marketing and advice on the acquisition and divestment of assets across various sectors. Previous to this role he spent five years as a stockbroker with a leading Australian-based stockbroking firm.

- **Adam Osseiran**: (**Non-Executive Director)** Dr. Osseiran joined the Board in September 2015. He has been involved with BRN since 2012, providing advice and assistance on several aspects of technology, applications and commercial opportunities. Dr. Osseiran is the co-founder and a director of Termite Monitoring and Protection Solutions Pty Ltd, founded in 2013, to exploit the unique Wireless Smart Probe acoustic termite detection technology, operating in the US$15B global pest control market. He is also Senior Technical Advisor to Mulpin (MRL) Ltd which has developed a new patented concept of embedding electronic components within a multi-layered printed circuit board. Dr. Osseiran holds a Ph.D. in microelectronics from the National Polytechnic Institute of Grenoble, France and an M.Sc. and B.Sc. from the University of Joseph Fourier in Grenoble. He is currently Associate Professor of Electrical Engineering at Edith Cowan University in Perth, Western Australia.

Figure 13: Board Remuneration

Board	Annual Remuneration (A$)
Peter van der Made	200,000
Mick Bolto	80,000
Neil Rinaldi	50,000
Adam Osseiran	50,000

Source: Company, Foster Stockbroking estimates.

Figure 14: Substantial Shareholders

Major shareholders	Interest
Peter van der Made	18.9%
Robert Mitro	13.9%
Anil Mankar	13.7%
Metals X Limited	7.8%

Source: IRESS, Foster Stockbroking.

FOSTER STOCKBROKING DIRECTORY

Name	Role	Phone	Email
Stuart Foster	Chief Executive Officer	+61 2 9993 8131	stuart.foster@fostock.com.au
Chris Francis	Executive Director	+61 2 9998 8167	chris.francis@fostock.com.au
Haris Khaliqi	Executive Director	+61 2 9993 8152	haris.khaliqi@fostock.com.au
Martin Carolan	Executive Director	+61 2 9993 8168	martin.carolan@fostock.com.au
Mark Fichera	Executive Director	+61 2 9993 8162	mark.fichera@fostock.com.au
Mark Hinsley	Executive Director	+61 2 9993 8166	mark.hinsley@fostock.com.au
Tolga Dokumcu	Execution & Dealing	+61 2 9993 8144	tolga.dokumcu@fostock.com.au
George Mourtzouhos	Execution & Dealing	+61 2 9993 8136	george.mourtzouhos@fostock.com.au

Foster Stockbroking Pty Ltd
A.B.N. 15 088 747 148 AFSL No. 223687
Level 25, 52 Martin Place, Sydney, NSW 2000 Australia
General: +612 9993 8111 Equities: +612 9993 8100 Fax: +612 9993 8181
Email: contact@fostock.com.au
PARTICIPANT OF ASX GROUP

Foster Stockbroking recommendation ratings: Buy = return >10%; Hold = return between –10% and 10%; Sell = return <-10%. Spec Buy = return > 20% for stock with very high risk. All other ratings are for stocks with low-to-high risk. Returns quoted are annual.

DISCLAIMER AND DISCLOSURE OF INTERESTS:

Foster Stockbroking Pty Limited (**Foster Stockbroking**) has prepared this report by way of general information. This document contains only general securities information or general financial product advice. The information contained in this report has been obtained from sources that were accurate at the time of issue. The information has not been independently verified. Foster Stockbroking does not warrant the accuracy or reliability of the information in this report. The report is current as of the date it has been published. In preparing the report, Foster Stockbroking did not take into account the specific investment objectives, financial situation or particular needs of any specific recipient. The report is published only for informational purposes and is not intended to be personal financial product advice. This report is not a solicitation or an offer to buy or sell any financial product. Foster Stockbroking is not aware whether a recipient intends to rely on this report and is not aware of how it will be used by the recipient. Before acting on this general financial product advice, you should consider the appropriateness of the advice having regard to your personal situation, investment objectives or needs. Recipients should not regard the report as a substitute for the exercise of their own judgment. The views expressed in this report are those of the analyst/s named on the cover page. No part of the compensation of the analyst is directly related to inclusion of specific recommendations or views in this report. The analyst/s receives compensation partly based on Foster Stockbroking revenues, including any investment banking and proprietary trading revenues, as well as performance measures such as accuracy and efficacy of both recommendations and research reports. Foster Stockbroking believes that the information contained in this document is correct and that any estimates, opinions, conclusions or recommendations are reasonably held or made at the time of its compilation in an honest and fair manner that is not compromised. However, no representation is made as to the accuracy, completeness or reliability of any estimates, opinions, conclusions or recommendations (which may change without notice) or other information contained in this report. To the maximum extent permitted by law, Foster Stockbroking disclaims all liability and responsibility for any direct or indirect loss that may be suffered by any recipient through relying on anything contained in or omitted from this report. Foster Stockbroking is under no obligation to update or keep current the information contained in this report and has no obligation to tell you when opinions or information in this report change.

Foster Stockbroking and its directors, officers and employees or clients may have or had interests in the financial products referred to in this report and may make purchases or sales in those the financial products as principal or agent at any time and may affect transactions which may not be consistent with the opinions, conclusions or recommendations set out in this report. Foster Stockbroking and its Associates may earn brokerage, fees or other benefits from financial products referred to in this report. Furthermore, Foster Stockbroking may have or have had a relationship with or may provide or has provided investment banking, capital markets and/or other financial services to the relevant issuer or holder of those financial products. For an overview of the research criteria and methodology adopted by Foster Stockbroking; the spread of research ratings; and disclosure of the cessation of particular stock coverage, refer to our website http://www.fostock.com.au.

Specific disclosure: The analyst has received assistance from BRN in preparing the report. The company has provided the analyst with communication with senior management and information on the company and industry). As part of due diligence, the analyst has independently and critically reviewed the assistance and information provided by BRN to form the opinions expressed in the report.

Specific disclosure: Foster Stockbroking is engaged by BRN to provide corporate services for which it has earned fees and continues to do so. Services include, but are not limited to, capital markets advisory, research, and marketing roadshows.

Specific disclosure: Foster Stockbroking is advising BRN on its 1-for-26 rights issue at $0.15/share announced April 2016 for which it is earning fees.

Specific disclosure: The analyst/s, Foster Stockbroking and/or associated parties may have (material) beneficial ownership in financial products issued by BRN at the time of this report. Diligent care has been taken by the analyst/s to maintain an honest and fair objectivity in writing the report and making the recommendation.