HOLISTIC POWER ANALYSIS OF IMPLEMENTATION ALTERNATIVES FOR A VERY LARGE SCALE SYNTHESIS ARRAY WITH PHASED ARRAY STATIONS

Andreea Anghel*Rik Jongerius†Gero Dittmann*Jonas Weiss*Ronald P. Luijten**IBM Research – Zurich, Switzerland†IBM Research, The Netherlandsaan@zurich.ibm.comr.jongerius@nl.ibm.com{ged,jwe,lui}@zurich.ibm.com

ABSTRACT

The Square Kilometre Array (SKA) will be the largest radio telescope in the world, generating data at Pb/s rates. Real-time processing will require 10^{18} compute operations per second and system operating costs will be dominated by energy consumption. In this paper we explore design options for the aperture array of the first SKA construction phase and provide lower bounds on their power consumption. We analyze the system's components from the antenna front-end to the central signal processor and identify the main power consumers. We compare ASIC-based and FPGA-based data processing pipelines and show that ASICs can lead to 1.6 to 4 times more power efficiency.

Index Terms— Square Kilometre Array, Design-Space Exploration, Power Consumption, FPGA and ASIC

1. INTRODUCTION

Big data challenges, such as processing large data volumes in realtime, prevent state-of-the-art radio telescopes from achieving the accuracy necessary to study radio signals that originated billions of years ago. The Square Kilometre Array (SKA) [1] is a nextgeneration telescope which aims to overcome these challenges. By providing an infrastructure that transports and processes data rates in the Pb/s range, SKA will be the largest and most precise radio telescope in the world. The aggregated data rate from all the antennas will be at least 2.5 Pb/s requiring as many as 10¹⁸ compute operations/second to process.

One of the main hurdles that the SKA system design will need to surmount is power consumption. In order to make the right choices early in the design process, we introduce a system model to estimate what the power envelope of a particular design will be for the first construction phase SKA1. Eight design points are analyzed and lower bounds of their power consumption presented. We investigate every subsystem of the signal processing chain and compare implementations with ASICs and FPGAs.

Our focus is the analysis of SKA1-Low, an aperture array of ~260,000 antennas, grouped in stations spread out over an area of 50 km in diameter, which operates in the 50 - 350 MHz frequency range. The radio signals follow a multi-stage processing chain: 1) they are filtered and amplified in the front-end; 2) the processed signals are sent to a station computing facility where the analog signals are digitized, Fourier transformed and beam-formed; 3) the station beams are sent for correlation to a central signal processor (CSP), located ~700 km away; 4) the correlated frequency channels are finally sent to a science data processor (SDP) for image processing.

2. RELATED WORK

Faulkner et al. [2] calculate the power budget of SKA1-Low assuming technology advances that are anticipated for 2016 and onwards. Our study complements this work with power models for each subsystem of SKA1-Low and power consumption estimates in both ASIC and FPGA technologies. Moreover, we provide details on how we scale the power consumption of the different subsystems.

D'Addario [3] proposes an ASIC-based concept for SKA and provides dynamic power consumption estimates for the SKA system. We extend this work by providing a higher level of detail in terms of power models and scaling rules. Moreover, we compare ASIC with FPGA implementation in 14 nm technology.

Jongerius [4] performs an analysis of the LOFAR telescope in terms of computational load and bandwidth requirements. We extend his models and estimate the power of the SKA-Low digital processing pipeline. Additionally, we provide models for signal transport and intra-station data motion over memory and interconnect.

Previous telescopes have mainly implemented their digital processing pipelines in FPGA technology [5, 6] and only very few in ASIC technology [7]. But, power consumption has not represented a design constraint for either of them. Thus, a holistic power analysis, as the one presented here, has not been essential until now, i.e. until the pre-construction phase of the SKA.

3. SYSTEM DESIGN POINTS

Table 1 gives an overview of the analyzed configurations. We compare two system design points: the baseline design [8] and an alternative design [9]. The baseline design features 911 stations, each covering 289 antennas and generating one beam. The alternative design features only 280 stations, but each covers 940 antennas and generates four beams. For both design points we analyze: 1) a single-stage versus a two-stage signal channelization and 2) single-stage digital beam-forming fully implemented in the station versus two-stage beam-forming implemented with analog in the front-end followed by digital in the station. By applying a polyphase filter (PPF), the singlestage channelization step implements a 1.14 kHz-channelization in the station only, while the two-stage channelization distributes the computational load between the stations and the CSP.

4. POWER MODELS

Fig. 1 shows in detail the SKA1-Low system model that we consider for power analysis. The number of RF streams (N'_{ant}) arriving from the antennas at the stations equals the number of antennas divided by the analog beam-former (ABF) tile size, i.e., one antenna (ABF not enabled) or four antennas (ABF enabled). In this paper we suppose that only one beam is generated per ABF tile. Moreover, we calculate the digital processing power primarily based on the energy consumption of a real-valued multiply-accumulate operation (MAC).

This work was conducted in the context of the ASTRON and IBM joint project, DOME, funded by the Netherlands Organisation for Scientific Research (NWO), the Dutch Ministry of EL&I, and the Province of Drenthe.

	Configuration	~		~	~ .	~-	~ ~	~-	~
	Parameter	Cl	C2	C3	C4	C5	C6	C 7	C8
STATION	No. antennas – N_{ant}	289	289	940	940	289	289	940	940
	Station diameter (m)	35	35	75	75	35	35	75	75
	ABF tile size (antennas)	1	1	1	1	4	4	4	4
	Subband bandwidth (KHz)	146.5	1.14	146.5	1.14	146.5	1.14	146.5	1.14
	No. subbands – $N_{subbands}$	2048	262144	2048	262144	2048	262144	2048	262144
	No. beams – $N_{stat-beams}$	1	1	4	4	1	1	4	4
CSP	No. stations – N_{stat}	911	911	280	280	911	911	280	280
	No. channels per subband – $N_{channels}$	128	1	128	1	128	1	128	1

Table 1. System configurations.

4.1. Power Models - Station

Front-end. The power model comprises the power consumed by the low-noise amplifiers (LNA) and, depending on the configuration, by the ABF. Given the relatively short distances between the antennas in a 4-element tile, we assume the power consumption of the ABF component to be negligible. Thus, provided the LNA power per receiver, the model is linear in N_{ant} and N_{pol} .

Antenna signal transport (to station). The RF signals need 300 MHz bandwidth and must be transported over a distance of 35 m to 75 m to the station processor. After the LNA, the signal power is ~ -70 dBm. At the ADC input, specifications require the level to be ~-45 dBm. As a consequence, a net link gain of > 20 dB is required. To compensate for cable losses, another ~10 dB is necessary, resulting in ~30 dB amplification required for the entire link. Assuming a 50/50 gain split between the TX and RX sides and one ADC per RF stream per polarization, the power consumption is linear in N'_{ant} .

Digital processing. The power consumption of the digitization, channelization and beam-forming steps is quantified based on their corresponding number of operations/second [4]. Given the energy necessary to generate one digital sample, the digitization power is linear in N'_{ant} and the ADC sampling rate ($f \ge 600$ MSamples/s, 8-bit samples). The PPF filter power includes the FIR filters and the FFT processing steps. An N_{taps} FIR requires N_{taps} MAC operations to filter a digital sample. We estimate the number of FFT operations based on the number of FFT butterfly operations required to process N input samples, i.e., $N_{butterfly} = \frac{N}{2} \cdot log_2(N)$, where $N = 2 \cdot N_{subbands}$. Eq. (1) and (2) show the PPF components power models, where $E_{R-butterfly}$ represents the energy per real-valued butterfly operation which consumes as much as ~1.5 MAC ops.

$$P_{FIR} = N'_{ant} \cdot N_{pol} \cdot (2 \cdot N_{subbands}) \cdot N_{taps} \cdot E_{MAC} \cdot \frac{f}{N} \quad (1)$$
$$P_{FFT} = N'_{ant} \cdot N_{pol} \cdot N_{butterfly} \cdot E_{R-butterfly} \cdot \frac{f}{N} \quad (2)$$

The digital beam-former (DBF) applies a phase delay on each frequency sample to steer the beam in a certain direction on the sky and then sums the resulting samples from all antennas per frequency subband. The phase delay is implemented as a matrix multiplication for all polarizations simultaneously and the summation per subband is implemented as an addition operation across all the RF streams. The DBF power model is detailed in Eq. (3), where E_{ACC} is the energy per real-valued accumulate/addition operation.

$$P_{DBF} = N_{subbands} \cdot N_{stat-beams} \cdot [N'_{ant} \cdot (2 \cdot N_{pol})^2 \cdot E_{MAC} + (N'_{ant} - 1) \cdot N_{pol} \cdot 2 \cdot E_{ACC}] \cdot \frac{f}{N}$$
(3)

Data motion to/from on-chip memory. The power required by this functional block includes both the dynamic and leakage power

necessary to move the input samples of an FFT operation to/from memory. Each length-N FFT operation consists of $log_2(N)$ stages of N/2 butterfly operations each. For every stage, N samples have to be available to start processing. Given the small amount of data (≤ 2 MB) we assume that the input samples are stored in on-chip memory. The power model is detailed in Eq. (4), where E_{sample}^{RAM} is the dynamic energy to read/write one sample from/to memory and P_{leak} is the memory leakage power.

$$P_{RAM} = N_{ant}^{'} \cdot N_{pol} \cdot \left(\frac{N}{2} \cdot log_2(N) \cdot 2 \cdot E_{sample}^{RAM} \cdot \frac{f}{N} + P_{leak}\right)$$
(4)

Chip-to-chip data I/O. ADC, PPF and DBF are assumed to be distributed over a three-chip architecture. We calculate the power needed to transport the total number of bytes per second between each pair of chips. Two chip-to-chip segments are identified: ADC-to-PPF and PPF-to-DBF. The relation between I/O power and data rate is assumed to be linear [10]. Therefore, the power of the ADC-to-PPF chip crossing is linear in N'_{ant} and ADC output data rate $(f \ge 4.8 \text{ Gb/s per ADC})$, while the PPF-to-DBF chip crossing power is linear in N'_{ant} , $N_{subbands}$ and FFT output data rate.

Data transport (from station to CSP). The station output is transported to the CSP over a distance of ~700 km at a data rate of up to 10 Gb/s (baseline design) or 40 Gb/s (alternative design). One suitable technology for this high bandwidth is to use optical transmission. Using dense wavelength division multiplexing (DWDM), the output of N'' stations can be transported on a single fiber. To overcome losses, the signal needs to be re-amplified at regular intervals by a repeater. Hence, the power is modeled as a fixed necessary TX/RX power at the end-points, plus an integer multiple of a fixed transport power, i.e., $P_{repeater,D}$ needed to regenerate the signal for a distance of D km.

4.2. Power Models – Central Signal Processor (CSP)

All N_{stat} stations send their output beams to the CSP which receives $M = N_{stat} \cdot N_{stat-beams} \cdot N_{subbands} \cdot N_{pol}$ subbands.

CSP digital processing. The CSP power is calculated based on the number of operations/second required for each of its processing steps [4]. The PPF power model is similar to the one in the station except that the input FIR/FFT samples are complex valued. The phase delay and the bandpass correction are implemented as a complex multiplication. Their power model is linear in M, $N_{channels}$ and the CSP PPF output data rate. The most power-intense CSP processing step is correlating all the subbands across each pair of stations $(N_{baselines} = \frac{N_{stat} \cdot (N_{stat} - 1)}{2})$, per polarization and per beam. The power model is detailed in Eq. (5), where E_{CMUL} is the energy consumption per complex-valued multiply operation.

$$P_{cor} = \frac{N_{baselines}}{N_{stat}} \cdot M \cdot N_{pol} \cdot E_{CMUL} \cdot \frac{f}{N \cdot N'}$$
(5)



Fig. 1. SKA1-Low model: (1) Front-end (LNA and analog beam-former) (2) Antenna signal transport from antenna to station (3) ADC (4) Chip-to-chip data I/O (5) Poly-phase filter (PPF = FIR + FFT) (5') On-chip memory (6) Chip-to-chip data I/O (7) Digital beam-former (DBF) (8) Data transport from station to CSP (9) CSP PPF (9') On-chip memory (10) CSP PD (phase delay) (11) CSP BC (bandpass correction) (12) Chip-to-chip data I/O (13) Correlation and integration. Parameters: N_{ant} – number of antennas per station, N_{pol} – number of polarizations per antenna, f – the ADC sampling rate, $N = 2 \cdot N_{subbands}$ – number of FFT points per (station) FFT block, $N_{stat-beams}$ – number of station output beams, $N' = N_{channels}$ – number of FFT points per (CSP) FFT block.

Subsequently, the output bandwidth is reduced by integrating the correlated data in time, by accumulating a certain number, e.g., 9600 of samples.

Data motion to/from on-chip memory. Analogous to the station processing, we assume the FFT input samples to be temporarily stored in on-chip memory. The dynamic and leakage power models are thus similar to the ones described in Sec. 4.1. The leakage power is linear in M, whereas the dynamic power is linear in both M and the output rate of the CSP FIR step.

Chip-to-chip data I/O. The PPF and the correlator (COR) are assumed to be distributed over a two-chip architecture. Thus, only one chip-to-chip interface is identified: PPF-COR. The I/O model is linear in M, $N_{channels}$ and the CSP FFT output data rate.

4.3. Power Parameters and Scaling Rules - Station and CSP

Front-end. We assume the LNA power to be ~30 mW per antenna and per polarization [2]. Recent studies already provide such LNA solutions customized to the SKA-Low requirements [11].

Antenna signal transport (to station). For this link we select copper cables because of their lower overall cost and sufficient performance. Our studies indicate that the cost-performance break-even between copper cables and fibers is reached at a link length of about 100 m, given 300 MHz bandwidth. Suitable 50 Ω line drivers with ~15 dB gain are available off the shelf and consume roughly 50 mW, adding a total of ~100 mW (TX and RX) to the total power-budget of each antenna [12].

Digitization. A recent research study demonstrates the availability of ADCs in 32 nm CMOS that consume only 3.1 mW at 8-bit and 1.2 GSamples/s [13]. This not only supports Nyquist sampling for SKA1-Low, but also for SKA instruments with wider and higher frequency band requirements.

ASIC vs. FPGA. We analyze the power consumption of the digital processing pipeline implemented in ASIC and FPGA. We estimate the dynamic power requirements when implemented in 90 nm technology and extrapolate the power consumption towards a 14 nm technology, which is expected to be available in 2016.

Synthesis tests (22 nm technology, 125 MHz clock) yield an ASIC MAC energy consumption of 9.6 pJ for 32-bit operands. This value is scaled to 90 nm and 14 nm technologies using the NMOS-FET dynamic power indicator values ($C \cdot V^2$) reported in the ITRS PIDS tables [14]. The scaling factors are 0.83 for 22 nm \rightarrow 14 nm and 1.86 for 22 nm \rightarrow 90 nm transition, where the MAC energy

scales quadratically with the bit width of the input operands. The same scaling factors are applied for the energy consumption of a single accumulation (ACC) operation, for which the baseline energy is 0.24 pJ in 22 nm. The ACC energy is scaled linearly with the bit width of the input operands.

We scale the energy consumption of FPGA MAC/ACC operations in 90 nm based on the FPGA vs. ASIC dynamic power measurements reported in [15]: 7.1x for FPGAs that use *hard-wired* building blocks (memories, multipliers, DSP) compared to ASICs. When moving to 14 nm technology, we consider this relative consumption factor to remain unchanged.

On-chip memory. For on-chip memories, we assume embedded DRAM for ASICs and SRAM for FPGAs. To estimate the station dynamic energy of a read/write memory operation and the memory leakage power, we use the CACTI tool [16] (the "pure RAM interface") with the following parameters: LP-DRAM for embedded DRAM cell, ITRS-HP for SRAM cell, 16 kB and 2 MB memory sizes for two-stage and single-stage channelization configurations. The newest technology supported by CACTI is 32 nm. Thus, for 14 nm technology, we first calculate the read/write dynamic energy for 32 nm using CACTI and then scale it to 14 nm using a factor of 0.8 [14]. As for the leakage power, the scaling factors to 14 nm technology are 2x and 1x, for SRAM and embedded DRAM, respectively [14]. The CACTI and scaled results are shown in Table 2. We estimate the CSP on-chip memory dynamic and leakage power using a similar approach, for an on-chip memory size of 512 bytes. For embedded DRAM and SRAM, respectively, this yields the following results: 0.32 pJ and 0.16 pJ dynamic energy consumption per 16-bit memory access and 21 μ W and 0.4 mW leakage power.

Table 2. 16-bit RAM access dynamic energy & leakage power.

Memory size	2 N	ИΒ	16 kB		
	32 nm	14 nm	32 nm	14 nm	
$DRAM_{dyn}$ (pJ)	51	40.8	5	4	
$SRAM_{dyn}$ (pJ)	40	32	2	1.6	
$DRAM_{leak}$ (mW)	14	14	0.23	0.23	
$SRAM_{leak}$ (mW)	900	1800	7	14	

Chip-to-chip data I/O. The power consumption for the off-chip interconnects depends on various factors such as data rate, chip-to-chip distance and channel requirements. To estimate the power consumption of the chip-to-chip transport, we use electrical interconnects, due to the relatively short distances involved. Current electri-



Fig. 2. a) Total power consumption b) CSP power consumption c) Station power distribution (ASIC) d) Station power distribution (FPGA). C1, C2, C5 and C6 refer to the baseline (Base) design, whereas C3, C4, C7 and C8 refer to the alternative (Alt.) design.

cal transceivers in 65 nm, capable of up to 15 Gb/s data rate, consume 75 mW [17] or 5 pJ/bit. For 14 nm we estimate an energy consumption reduction of 30%, which yields an energy of ~3.5 pJ/bit.

Transport to CSP. To transport data from the stations to the CSP (700 km), we select optical single mode transmission, because of their high data rates and long distances capabilities. Early studies [18] show that off-the-shelf DWDM repeaters are available and able to transport up to 160 Gb/s data over ~130 km of optical fiber, dissipating 25 W ($P_{repeater,130}$). To multiplex and demultiplex multiple channels on one optical fiber, two TX and RX modules are needed requiring 25 W each [19].

5. RESULTS AND DISCUSSION

Fig. 2 (a) shows the distribution of the estimated power consumption for the entire SKA1-Low system across all eight configurations in 14 nm ASIC and FPGA technologies. System power consumption ranges from 45 kW to 380 kW for ASIC and from 100 kW to 1500 kW for FPGA implementations. In both cases the two power contributors are the stations and the CSP. Fig. 2 (b) illustrates the CSP contributor for each configuration and shows that it only accounts for a small portion (~3-7% on average) of the total power. Indeed, the CSP consumes between only 1.65 and 3.45 kW in ASIC and between 8.80 and 24 kW in FPGA. Therefore, we will focus on the dominating power contributor (the stations), which is illustrated in Fig. 2 (c) and (d), in detail with its main components.

Each station consumes between 51 W and 1300 W in ASIC and between 86 W and 5300 W in FPGA technology. As shown in the charts, the main reason for the discrepancy between ASIC and FPGA implementations (~1.6 to 4 times in favor of the ASIC solution) is the much higher efficiency of the DBF, PPF and on-chip RAM components when implemented in ASIC. Fig. 2 (d) illustrates that the three alone account for the entire difference.

Focusing on the ASIC implementation, Fig. 2 (a) and (c) show that the alternative designs (configurations C3, C4 and C7, C8) exhibit a significantly larger per-station power consumption (between 2.7x and 3.5x) than their baseline counterparts. However, the total system power consumption is similar. Thus, the baseline and the alternative designs only expose a choice between performing the computation more centrally (in the CSP) or more distributively (in the stations), without significant impact on the total power requirements.

The four ABF-enabled configurations (C5, C6, C7, C8), through the reduction in both computation and data transfer, are more energy efficient in all cases. However, this comes at the price of reduced telescope flexibility, e.g., in terms of weighting / tapering of individual antenna signals. Therefore, enabling ABF exposes an efficiency/capability trade-off. For the currently targeted telescope capabilities, we will focus further on DBF configurations.

The last design choice we analyze is whether to use a single or dual-stage PPF. When the CSP requirements are deemed too high, resorting to a single-stage PPF (C2, C4, C6, C8) allows reducing the CSP load by shifting part of the computation towards the stations.

Finally, the power distribution charts clearly show that data motion plays the more significant role (compared to computations on data) in the total power consumption of the system. For example, in ASIC technology, off-station data transport can account for as much as 32% of the total station power (22% on average). Furthermore, if we exclude the front-end (which accounts for 16% on average), intra-station data motion over memory and between chips can reach up to 85% of the remaining station power (73% on average).

6. CONCLUSIONS AND FUTURE WORK

We presented a holistic SKA1-Low power analysis and provided lower power bounds for eight system configurations. The results show that an ASIC-based SKA implementation consumes between 1.6 and 4 times less total power than an FPGA-based implementation. Our analysis enables system designers to reason on the distribution of resources between station processing, data transport and correlator, and their subsystems. The high cost of intra-station data motion over memory and chip-to-chip data I/O requires a softwarehardware design approach which reduces data communication. The models detailed in this work will enable the analysis of the second phase of SKA and allow us to derive conclusions on the power consumption of the entire telescope.

7. ACKNOWLEDGMENTS

We wish to thank Peter Buchmann, Ton Engbersen, Martin Schmatz, Christoph Hagleitner, Jan van Lunteren, Silvio Dragone, Thomas Toifl, Stefan Wijnholds and Albert-Jan Boonstra for their support.

8. REFERENCES

- C. Broekema, A.-J. Boonstra, V. Cabezas, et al., "DOME: Towards the ASTRON & IBM center for exascale technology," in *Proceedings of the 2012 Workshop on High-Performance Computing for Astronomy Date*, June 2012, pp. 1–4.
- [2] A. Faulkner, D. Kant, P. Alexander, et al., "The aperture arrays for the SKA - the SKADS white paper," April 2010, Memo 122.
- [3] L.R. D'Addario, "How to implement SKA digital signal processing so that it uses very little power," in Workshop on Power Challenges of Mega-Science: The Example of SKA, June 2012.
- [4] R. Jongerius et al., "LOFAR retrospective analysis and application to SKA phase 1 and phase 2," Tech. Rep., IBM Research, to appear.
- [5] M.P. Van Haarlem, M.W. Wise, A. Gunst, et al., "LOFAR: The low frequency array," in *Instrumentation and Methods for Astrophysics*. May 2012, arXiv.
- [6] C. D. Patterson, B. S. Martin, S. W. Ellingson, et al., "FPGA cluster computing in the ETA radio telescope," in *IEEE International Conference on Field Programmable Technology*, December 2007, pp. 25–32.
- [7] A. Baudry, R. Lacassec, R. Escoffierc, et al., "Performance highlights of the ALMA correlators," in *SPIE 8452, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, July 2012.
- [8] SKA Organisation, "The Square Kilometre Array," http: //www.skatelescope.org.
- [9] S. Wijnholds and R. Jongerius, "Computing cost of sensitivity and survey speed for aperture array and phased array feed systems," in *IEEE Africon*, Mauritius, September 2013.
- [10] S. Palermo, A. Emami-Neyestanak, and M. Horowitz, "A 90 nm CMOS 16 Gb/s transceiver for optical interconnects,"

in *IEEE Journal of Solid-State Circuits*, May 2008, vol. 43, pp. 1235–1246.

- [11] M. Panahi, LNA Considerations for Square Kilometre Array, Ph.D. thesis, University of Manchester, 2012.
- [12] Mini-Circuits, "Amplifiers," http://www. minicircuits.com/products/Amplifiers. shtml.
- [13] L. Kull, T. Toifl, M. Schmatz, et al., "A 3.1 mW 8b 1.2 GS/s single-channel asynchronous SAR ADC with alternate comparators for enhanced speed in 32 nm digital SOI CMOS," in *IEEE International Solid-State Circuits Conference*, February 2013, pp. 468–469.
- [14] International Roadmap Committee, "International Technology Roadmap for Semiconductors 2012 Update," http://www. itrs.net.
- [15] I. Kuon and J. Rose, Quantifying and Exploring the Gap Between FPGAs and ASICs, Springer US, 2010.
- [16] HP Labs, "CACTI," http://www.hpl.hp.com/ research/cacti/.
- [17] G. Balamurugan, J. Kennedy, G. Banerjee, et al., "A scalable 5-15 Gbps, 14-75 mW low-power I/O transceiver in 65-nm CMOS," in *IEEE Journal Solid-State Circuits*, 2008, vol. 43, pp. 1010–1019.
- [18] JDSU, "L-band multichannel erbium-doped fiber amplifier (EDFA) WRA-217L," http://www.jdsu.com/ productliterature/wra217L_ds_cms_tm_ae. pdf.
- [19] JDSU, "DWDM transponder with multi-rate XFP interfaces," http://www.jdsu.com/productliterature/ wrt840_ds_cms_tm_ae.pdf.